



**10<sup>th</sup> International  
Conference on User  
Modeling**

**Edinburgh, UK  
24<sup>th</sup>-29<sup>th</sup> July 2005**

**<http://www.um.org/>**

**Workshop 9**

**PRIVACY-ENHANCED PERSONALIZATION  
(PEP2005)**

**Proceedings**

**Edited by: Alfred Kobsa and Lorrie Cranor**

**<http://www.isr.uci.edu/pep05/>**



***PEP2005***

***Edinburgh, Scotland***

**Proceedings of the UM05 Workshop on  
Privacy-Enhanced Personalization**

**Editors: Alfred Kobsa**

**Lorrie Cranor**

**July 25, 2005**



## Preface

Personalizing people's interaction with computer systems entails gathering considerable amounts of data about them. As numerous recent surveys have consistently demonstrated, computer users are very concerned about their privacy. Moreover, the collection of personal data is also subject to legal regulations in many countries and states. Both restrictions impact a number of frequently employed personalization methods. The aim of this workshop was to explore the potential of research on "privacy-enhanced personalization," which aims at reconciling the goals and methods of user modeling and personalization with privacy constraints imposed by individual preferences, conventions and laws.

Workshop participants were asked to consider, e.g., the following questions:

- How much personal data do individual personalization methods really need? Can we find out in advance or in hindsight what types of data contribute to reasonably successful personalization in a specific application domain, and restrict data collection to these types of data?
- Is client-side personalization a possible answer to privacy concerns and legal restrictions? What technical, legal and business obstacles will have to be overcome?
- In what way should the user be involved in privacy decisions? What does appropriate notice and choice look like, and what rights must and should be granted?
- Will we need trusted third parties, and what services will we need them to provide?
- How much can we benefit from anonymity or pseudonymity infrastructures, and are there limits that should be observed?
- Are distributed user models an answer or a problem from a privacy perspective?
- Does personalization in a mobile context pose additional challenges? How can they be overcome?
- Do mobile user models pose additional privacy problems?
- How can multi-user personalized systems cater to the privacy constraints of each individual user?
- What should an ideal legal framework look like from the perspective of privacy-enhanced personalization?
- Are special provisions necessary in the case of people with disabilities and student-adaptive educational systems?

Sixteen submissions were received, and about 1/4 were accepted for "long" presentation and 1/3 for a "short" presentation. The conference program also included an invited talk by Caspar Bowden, Microsoft's Chief Privacy Advisor for Europe, Middle East and Asia, and two discussion sessions.

The workshop benefited considerable from expertise of its program committee members whose assistance in the selection of papers was invaluable:

John Canny, University of California, Berkeley, CA

Clare-Marie Karat, IBM Watson Research Center, Hawthorne, NJ

Judy Kay, University of Sydney, Australia

Sarah Spiekermann, Humboldt University, Berlin, Germany

Loren Terveen, University of Minnesota, Minneapolis, MN

Last but not least, thanks are also due to Yang Wang for setting up the workshop website and for preparing the proceedings, as well as to UC Irvine's Institute for Software Research for its administrative support.

Alfred Kobsa and Lorrie Cranor

Workshop Co-Chairs

# **PEP05 Program**

***Monday, July 25***

8:30-9:15 INTRODUCTION

Organizational remarks

Introduction of participants

P Where Personalization, Privacy, and Security Meet (Position Statement)

Chris C. Demchak and Kurt D. Fenstermacher

University of Arizona, United States..... 1

9:15-10:15 INVITED INDUSTRY TALK

The Need for an Identity Meta-System

Caspar Bowden

Microsoft Chief Privacy Advisor for Europe, Middle East and Asia

10:15-10:45 COFFEE BREAK

10:45-12:30 USER STUDIES

L Perceived Control: Scales for Privacy in Ubiquitous Computing Environments

Sarah Spiekermann

Humboldt University Berlin, Germany..... 3

S Privacy & Personalization: Preliminary Results of an Empirical Study of Disclosure Behavior

Evelien Perik, Panos Markopoulos, Eindhoven University of Technology, The Netherlands

Boris de Ruyter, Philips Research Eindhoven, The Netherlands..... 15

S Informed Consent to Address Trust, Control, and Privacy Concerns in User Profiling

Thea van der Geest, Willem Pieterse, and Peter de Vries

University of Twente, The Netherlands..... 23

Discussion: Where to go next in user studies on privacy-enhanced personalization?

12:30-1:30 LUNCH

## 1:30-2:23 CATERING TO PRIVACY REQUIREMENTS

- L A Software Product Line Approach for Handling Privacy Constraints in Web Personalization  
Yang Wang and Alfred Kobsa  
University of California, Irvine, United States.....35
- S Privacy and Security in Ubiquitous Personalized Applications  
Ajay Brar, Judy Kay  
University of Sydney, Australia.....47

## 2:23- 3:15 SELECTIVE ACCESS TO USER DATA

- L A Single Sign-On Identity Management System Without a Trusted Third Party  
Brian Richardson and Jim Greer  
University of Saskatchewan, Canada.....55
- S Intra-Application Partitioning of Personal Data  
Katrin Borcea, Hilko Donker, Elke Franz, Katja Liesebach, Andreas Pfitzmann, and Hagen  
Wahrig  
Dresden University of Technology, Germany.....67

## 3:15-3:45 COFFEE BREAK

## 3:45 - 4:40 PRIVACY IN RECOMMENDER SYSTEMS

- L Privacy-Enhanced Collaborative Filtering  
Shlomo Berkovsky, Yaniv Eytani and Tsvi Kuflik<sup>1</sup>, University of Haifa, Israel  
Francesco Ricci, ITC-irst, Trento, Italy.....75
- S Privacy, Shilling, and The Value of Information in Recommender Systems  
Shyong K Lam and John Riedl  
University of Minnesota, United States.....85

## 4:40-5:30 DISCUSSION: WHERE TO GO NEXT IN PRIVACY-ENHANCED PERSONALIZATION?

AUTHOR INDEX.....93

# (Position paper) Where personalization, privacy, and security meet

Chris C. Demchak<sup>1</sup> and Kurt D. Fenstermacher<sup>2</sup>

<sup>1</sup> School of Public Administration and Policy and <sup>2</sup> Management Information Systems, Eller College of Management, University of Arizona, Tucson, AZ, 85721, United States  
<sup>1</sup> [demchak@u.arizona.edu](mailto:demchak@u.arizona.edu) and <sup>2</sup> [kurtf@eller.arizona.edu](mailto:kurtf@eller.arizona.edu)

**Abstract.** We have been developing dynamic user modeling techniques, while also pursuing policy research to strike a balance between an individual's privacy and society's security. We analyze user modeling through our policy lens, known as the behavior-identity knowledge (BIK) framework and offer suggestions on how to protect user privacy.

Existing work by Kobsa [1] and Cranor [2] has highlighted personalization's risks to privacy — to personalize systems requires gathering personal data, which is then used to guide the adaptation process. Much of this personalization can be captured by the single question, "What will the user do next?" By anticipating the answer, systems can better serve users by adapting the presentation of information [3] and other user interaction aspects. We this conflict between personalization and privacy as similar to national security concerns in a post-9/11 world. Rather than asking "What will the user do next?", however, people ask, "What will the suspect do next?" Instead of gathering data on user preferences, new profiling and tracking technologies accumulate data on suspects and others. Indeed, even data gathered in service of user modeling might later be used to hunt for terrorists and others. Advances in scale, scope, and the accuracy of user modeling inevitably place these technologies squarely in the debate on how to balance security with freedom, and particularly the freedom of privacy.

In previous work on balancing privacy and security, we have described privacy as an aggregate of two independent concepts (shown in Fig. 1): knowledge of behavior and knowledge of identity, which we call the Behavior-Identity Knowledge (BIK) model. We argue that privacy is not at risk unless an organization (or a person) knows both a person's identity and behavior. From a policy perspective compromising between knowing one or other, society should focus its efforts on monitoring behavior, initially without regard to identity. Only when there is reasonable cause, do we allow the institution revelation of the identity of suspicious persons. Moreover, we must have BIK-implementing institutional safeguards such that, whenever organizations consider behavior and identity together, we can quickly

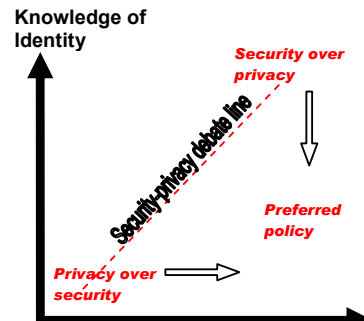


Fig. 1. Decomposing behavior and identity

validate the underlying data and offer a rapid appeals process to redress errors in the systems.

We will start small this fall with a simulation of this interplay using randomly created identities with varying identifiers including simulated fingerprints and DNA, and personal goals. We will use AI planners with time-stamped action sequences, and unrelated human like random actions. Simulated institutions will mask identities with validation and appeal (V&A) processes and also monitor stochastically the actions to model the difficulty in learning everything about everyone. Including organizational sharing of information as well, we will work to create a modular design that would enable varied user modeling techniques to apply in the simulation. We will model the organizations as searching for suspicious patterns of behavior, but the agencies will not have access to the plans and goals of the population, but instead only the actions that are the realization of those plans and goals. In the simulation, organizations will be able to petition for the resolution of a pseudonym once the likelihood of a match between a person's actions and a suspicious pattern of activity exceeds a threshold, just as law enforcement agencies must today meet escalating burdens to authorize more invasive actions against citizens.

In this work, we address the challenge of security, privacy, and the double-edged sword of advancing personalization in a widely networked society. We suggest a framework and hope to provide a design to help resolve this dilemma technically and institutionally, in the form of a simulation to test both the BIK framework and user modeling techniques in a controlled environment.

## References

- [1] A. Kobsa and J. Schreck, "Privacy through pseudonymity in user-adaptive systems," in *ACM Transactions on Internet Technology*, vol. 3, pp. 149-183, 2003.
- [2] L. F. Cranor, "'I didn't buy it for myself' privacy and ecommerce personalization," presented at 2003 ACM workshop on privacy in the electronic society, Washington, D. C., U.S.A., 2003.
- [3] A. Kobsa, "Personalized hypermedia and international privacy," in *Communications of the ACM*, vol. 45, pp. 64-67, 2002.
- [4] E. Alderman and C. Kennedy, *The Right to Privacy*. New York, NY, USA: Alfred A. Knopf, 1995.
- [5] R. O'Harrow, Jr., "In Age of Security, Firm Mines Wealth Of Personal Data," in *The Washington Post*. Washington, D.C., 2005, pp. A1.
- [6] C. C. Demchak and K. D. Fenstermacher, "Balancing security and privacy in the information and terrorism age: distinguishing behavior from identity institutionally and technologically," in *The Forum*, vol. 2, pp. Article 6, 2004.

# Perceived Control: Scales for Privacy in Ubiquitous Computing

Sarah Spiekermann<sup>1</sup>

Humboldt University Berlin, Institute of Information Systems,  
Spandauer Strasse 1, 14057 Berlin, Germany  
[sspiek@wiwi.hu-berlin.de](mailto:sspiek@wiwi.hu-berlin.de)

**Abstract.** Ubiquitous computing (UC) environments have triggered a strong research interest in privacy. How can people remain private when the infrastructure and objects around them begin to talk? Heading for an answer to this question many studies have rushed over past years to present guidelines for privacy-friendly UC design and have tempted even to rewrite the vocabulary of this socio-psychological construct. In doing so, most authors notice though that when it comes to requirements specification for privacy in UC, user-friendly technology design is really more about *perceived control* than it actually is about the end state of privacy itself. The current position statement therefore attempts to pull the two constructs –privacy and control- apart by theoretically reflecting on their mutual dependencies. It then proceeds by proposing a scale for appropriate measurement of perceived control in UC environments.

## 1 Introduction

Privacy is a construct widely investigated in the Information Systems world, both in the context of E-Commerce as well as in the context of UC. Except for a few articles what has been missing in IS research though is a thorough framing and defining of what privacy is including empirical testing of its building blocks based on properly defined scales. As a result of this lack of research, privacy definitions appear in different forms and facets, misconceptions not excluded. Consequently, when researching privacy for a Ubiquitous Computing context today, there is little common ground to build on.

Ubiquitous Computing refers to environments where most physical objects are enhanced with digital qualities. It implies “tiny, wirelessly interconnected computers that are embedded almost invisibly into just about any kind of everyday object” [1]. Thus, people buy and use products that can be automatically recognized, tracked, addressed and, potentially, trigger activities or services. Because of these properties, UC and especially one of its core technologies, RFID, have stirred some strong debates about privacy being at risk.<sup>1</sup>

---

<sup>1</sup> RFID chips (tags) are embedded into the fabric of products and emit a unique product number once addressed by a reader. The reader feeds the number in a backend information

Yet, a misconception of privacy is actually articulated already in one of the most widely cited articles on Ubiquitous Computing, notably Mark Weiser's "The Computer for the 21<sup>st</sup> Century" [2]. Commenting on social challenges arising from UC, Weiser wrote: „The [social] problem [associated with UC], while often couched in terms of privacy, is really one of control." While Mark Weiser was right to point out that UC raises control issues reaching beyond privacy alone, it should be noted is that privacy has actually for decades been defined in terms of control. Altman [3] for example, one of the founding fathers of privacy research in the Western hemisphere, defined privacy in 1975 as "the selective control of access to the self or to one's group." Schoeman [4] saw privacy as "the control an individual has over information about himself or herself." And Margulis [5] reflected on several decades of privacy research when writing: "Privacy, as a whole or in part, represents control over transactions between person(s) and other(s), the ultimate aim of which is to enhance autonomy and/or minimize vulnerability." Summing up, privacy cannot be seen as separate from control. Instead it is deeply intertwined with it.

Unfortunately, IS research has seen few works building upon this fundamental insight.. For this reason, we want to investigate privacy more systematically with a view to its inherent control character. More specifically, we want to develop scales that are able to measure *perceived* privacy governance on the basis of perceived control. UC serves as the context in which privacy is sought.

## 2 Loss of Privacy through Loss of Control in UC

Loss of privacy in UC environments can really be due to two distinct reasons: The first one is relating to what we want to call '*people losing control over being accessed*'. In classical privacy literature, researchers relate to this aspect of privacy when they discuss the *collection* of data by marketers and other institutions [6]. For UC environments it is typically assumed that sensor- and RFID infrastructures will be ubiquitous. The so called "intelligent infrastructure" seeks to automatically adapt to people moving through space and for this it needs to establish connections with peoples' objects. People are envisaged to be read out by RFID readers or be tracked by other technologies. Building on Altman [3], Boyle refers to this privacy aspect in UC as the need "to control the *attention* of the Ubicomp environments" [7]. This control can be exercised through Privacy Enhancing Technologies (PETs). PETs – according to current research – are supposed to enable users to protect themselves from being accessed against their will. First PETs for UC are blocker tags [8], the Privacy Awareness System (pawS) [9] or authentication based protection schemes [10-12].

The second factor impacting privacy in UC is due to a *lack of control over information use and maintenance* once people (or their objects) have been accessed. This concern about unauthorized secondary use is actually a historical one in privacy research [6]. However, UC adds a new dimension of relevance to this aspect of privacy since much more data is being collected. Ubiquitous multimedia

---

infrastructure where the nature of the product and potentially its owner is identified. Based on this information, further services are being triggered.

environments can, for example, lead to a more prevalent risk of disembodiment or disassociation as discussed by Belotti and Sellen [13]. Tracking of whereabouts and social network analysis suddenly gain a ‘physical’ dimension [14]. And, unique item identification inherent in new numbering standards, such as the Electronic Product Code (EPC) or IPv6 can lead to a degree of personal attribution and potential surveillance unseen before.

Still, this secondary use (and abuse) of information is not possible if there has not been access in the first place. This implies that controlling access is a crucial part of the privacy equation in UC. We therefore focus on the first dimension of privacy in UC: perceived control over the *access* that intelligent infrastructures may gain to individuals via their objects.

We proceed as follows: In section 3 we introduce the reader to the main theories of perceived control mostly deduced from the past 30 years of psychological research. Furthermore, we comment on two main privacy enhancing technologies envisioned to induce control over UC in people. Based on this, we then describe the development of scales that are able to measure control over UC perceived by PET users (section 4). We then apply these scales to a UC

### **3 Perceived Control and PETs in UC Environments**

#### **3.1. Perceived Control**

Perceived control is a construct investigated in psychology since the 1960s [15]. One of the first investigations of control can be found in Seligman’s work on *learned helplessness* [16]. Learned helplessness was considered by Seligman as the opposite of being in control. Together with Abramson et al. [17] he defined helplessness as “cases in which the individual ... does not possess controlling responses” (p.51). People enter into a stage of numbness where they feel that their activity really does not impact the course of activities around them. In the context of UC this would imply that people have given up on protecting their privacy as they believe protection efforts to be in vain anyways.

Related to this feeling, but somewhat weaker in emotional strength is the notion of control as a means to achieve a desired outcome. Seligman propagated this aspect noting that “a person has control when a desired outcome’s occurrence is dependent on the person’s responses” [18]. In psychological research this position has mostly been referred to as *contingency* [19].

While Seligman and his peers’ research focused on response contingency, Langer propagated that people can only perceive control if they are aware that they can influence these through their *choices*: “...control...is the active belief that one has a choice among responses that are differentially effective in achieving the desired outcome” [18]. In a UC environment this choice aspect would imply that people can opt easily out of being accessed by the intelligent infrastructure.

In order to recognize one’s choices a major requirement is that one is properly informed about one’s options. As Fiske and Taylor [20] put it: “...a sense of control

...is achieved when the self obtains or is provided with information about a noxious event” (p.201). Skinner calls this type of control “information control” [15]. In a UC context, *information control* would mean that people are not read out without them being aware of it.

Moreover, there is a *power* aspect in control that has been considered by Rodin writing: “[perceived control is]...the expectation of having the power to participate in making decisions in order to obtain desirable consequences and a sense of personal competence in a given situation” [21]. In fact, power is an important notion also in the literature on motivation. When people feel power they may be motivated to use a technology more rigorously.

Yet, Rodin also referred here to another notion of control which is one’s feeling of competence. If people do not feel competent enough to master a situation, they will not feel in control. Bandura is one of the scholars focusing on this aspect of control which is referred to as self-efficacy: “people’s beliefs about their capabilities to exercise control over events that affect their lives” [22]. Researchers in technology acceptance also use the term ‘*ease-of-use*’ in this context [23]. They see control as an important antecedent for peoples’ impression (or experience) to easily handle a new technology [24].

### 3.2. PETs for Perceived Control in UC

The goal of this article is to document the development of scales that are able to measure to what degree UC privacy enhancing technologies (PETs) are able to induce a perception of control in people. We assume that if people perceive control over UC environments through their PETs then they will also perceive themselves exercising their right to privacy. Before delving into the details of scale development yet it is important to describe available PETs for UC in more detail and to give the reader a perspective on the type of PETs used to test the control scales reported on hereafter.

Based on current UC PET research, we consider two types of privacy enhancing technologies as important in the context of RFID technology. We term these two alternative PET approaches the ‘user model’ and the ‘agent model’.

The *user model* implies that users exert full control over RFID tags by means of appropriate authentication mechanisms. Objects do not a priori respond to network requests. Instead the user self-initiates the use of intelligent services if they are available and useful in the respective context. The context decision when and how the use of tags is appropriate is thus taken by the object owner [10-12, 25]. If the owner of the object has some benefit from reviving an object’s RFID tag she can do so by authenticating access using a password. We expect the user model to induce a high level of control with users since the intelligent infrastructure cannot act autonomously.

In contrast, the *agent model* is based on the idea that RFID tags are left on active by default, thus always answering to network requests. Access control in this scenario is provided automatically via consumer privacy preferences that are residing in the network and are exchanged via some identity management system. This system takes the context-decision for the object owner when to answer network requests and when to deny them. When the network seeks access to a person (e.g. in order to send an

advertising message or track a person's movements), an identity management system, or agent, matches personal privacy preferences with claimed data collection purpose [9, 26]. Here, it is the an external party communicating with the agent that typically initiates communication. The user does not have to become active. The user trusts that his agent and the network interacting with it adhere to his privacy preferences.

## **4 Scale development and testing**

### **4.1. Control Definition and Initial Item Development**

Based on the control literature described in section 3.1. we developed scales that would be able to test peoples' perceived control over being accessed by an intelligent infrastructure. Helplessness, contingency, choice, power, information and ease-of-use described above served as the basic categories to frame the construct (see table 1).

Following the guidelines of proper scale development [27], the first step was the development of a proper definition of the perceived control construct. Based on an expert we formulated the following definition: "*Perceived control [in a UC environment] is the belief of a person in the electronic environment acting only in such ways as explicitly allowed for by the individual.*" We then developed 14 questions (items) capturing the different control categories identified above. To assess the relatedness of these items with the control construct definition we then conducted interviews with 25 participants (mostly students). Participants ranked the 14 questions in an order of decreasing relatedness to the control definition. Ten participants furthermore categorized the items into meaningful categories that matched the different control aspects we had hoped to capture. Based on this ranking and classifying we were able to identify three questions that were the least related to the definition and we excluded them from further research. In parallel we adjusted four questions from the Technology Acceptance Model on ease-of-use to fit our context [23, 24]. The resulting 15 questions promised a high degree of content validity. Their importance ranking with regards to the control definition and their respective categories are presented in table 1. The next step was to test whether these categories would indeed show and be internally consistent when applied to UC PETs.

### **4.2. Empirical Item Testing**

128 subjects were invited by a market research agency to participate in a study on tomorrow's shopping environments. They were demographically representative with 47% female and 53% male. 36% were below 30 years of age, 21% 30 to 39 and 43% 40 years or older. 40% had no A-levels and only 25% went to university. 81% had an income below € 30.000.

The participants were split into two random groups. Group 1 contained 74 subjects. Group 2 had 54 participants. Both groups were presented with a film on future shopping environments in which RFID technology would be used. RFID technology,

representing the UC environment here, was explained neutrally. Its benefits and drawbacks were commented on without bias. After-sales benefits of RFID were described on the basis of two services: an intelligent fridge and product return without need for a receipt. The film was identical for both groups except for one piece of information: the privacy enhancing technology (the PET) available to the consumer to control his privacy. In group 1 the film briefing was such that RFID chips would all be switched off at the supermarket exit but could be turned on again with the help of a personal password if after-sales services (fridge, product exchange) would require so (user model). In group 2 the film briefing was such that chips would all be left on at the supermarket exit but could only be accessed by readers for after-sales purposes if the reading purpose would match a person's privacy preference (agent model).

Before and after seeing the film participants answered a battery of questions. The 15 control items were passed among other questions after the film. As depicted in table 1 they were answered on a 5-point Rohrman scale [28].

**Table 1.** Control items and categories

Rank	Index	Question text <i>(1 = fully agree ... 5 = do not agree at all)</i>	Category
1	POW 1	I feel that I can steer the intelligent environment in a way I feel is right.	Power
2	POW 2	Thanks to <the PET> the electronic environment and its reading devices will have to subdue to my will.	
5	POW 3	Due to <the PET> I perceive perfect control over the activity of my chips.	
3	CON 1	Thanks to <the PET> I could determine myself whether or not I'll interact with the intelligent environment.	Contingency
7	CON 2	Through <the PET>, services are put at my disposition when I want them.	
6	H 2	I could imagine that if the electronic environment set out to scan me, it would be able to do so despite <the PET>.	Helplessness
10	H 1	<The PET> will finally not be able to effectively protect me from being read by the electronic environment.	
8	COI 1	Due to <the PET> it is still my decision whether or not the intelligent environment recognizes me.	Choice
4	COI 2	Through <the PET> I finally have the choice whether or not I am being scanned or not	
9	IC 1	Through <the PET> I would always be informed of whether and in what form the electronic environment recognizes me.	Information
11	IC 2	Using <the PET> I would always know when and by whom I have been read out.	
*	EUP 1	To learn to use <the PET> would be easy for me.	Ease-of-use
*	EUP 2	It would be easy for me to learn skillful use of <the PET>.	
*	EUP 3	I would find <the PET> easy to use.	
*	EUP 4	Due to <the PET> the information exchange between my chips and reading devices would be clearly defined.	

### 4.3. Internal Consistency and Reliability of Control Items

To understand whether the six control categories would really be reflected in the 15 control related questions we first conducted factor analysis. Assuming that there could be correlations between factors we chose oblimin rotation. Very few missing items were replaced by mean values. Principal component analysis was employed. Factor analysis was first conducted for group 1 (user model) and it was then analysed whether the results would replicate for group 2 (agent model). This first round of analysis showed that only 8 out of the 15 questions consistently (across both treatments) load on three factors with factor loadings above .6. 2 items, one ease-of-use question and one question on contingency saw low loadings for both treatments and were therefore eliminated from the item set. 5 remaining questions, notably those on power and choice would not load consistently on the three factors. In fact, for group 1 power and choice related questions loaded together with information and contingency items. Group 2 saw power and choice loading with helplessness. We therefore concluded that the items developed for power and choice would not be suited to reliably distinguish between factors and we opted to eliminate them from the list of questions, well recognizing that content validity of left over scales would suffer due to this step. The remaining 8 questions were used again to first run factor analysis for group1 and then (to confirm reliability) for group 2. In this step, three factors explaining the perceived control construct could clearly be identified for both PET samples (see table 2).

**Table 2.** Final factor loadings for the 2 PET treatments

Password PET (group 1)				Agent PET (group 2)			
	Pattern Matrix(a)				Pattern Matrix(a)		
	1	2	3		1	2	3
EUP 2	<b>0,954</b>	-0,048	-0,021	EUP 2	<b>0,937</b>	0,042	-0,028
EUP 1	<b>0,881</b>	-0,065	-0,094	EUP 1	<b>0,925</b>	-0,056	-0,045
EUP 3	<b>0,854</b>	0,162	0,088	EUP 3	<b>0,905</b>	0,047	0,074
IC 2	-0,114	<b>0,918</b>	-0,046	IC 2	-0,069	<b>0,880</b>	-0,024
IC 1	0,077	<b>0,855</b>	0,067	IC 1	0,026	<b>0,872</b>	0,004
CON 1	0,068	<b>0,822</b>	-0,025	CON 1	0,082	<b>0,847</b>	0,023
H 2	0,109	-0,014	<b>0,905</b>	H 2	0,062	-0,159	<b>0,877</b>
H 1	-0,165	0,001	<b>0,800</b>	H 1	-0,068	0,180	<b>0,801</b>

Rotation Method: Oblimin with Kaiser Normalization.  
a. Rotation converged in 5 iterations.

Rotation Method: Oblimin with Kaiser Normalization.  
a. Rotation converged in 4 iterations.

Factor 1 is clearly related to the category ‘ease-of-use’ of the PET. The three questions (EUP 1, 2, 3) measure to what extent one feels control over RFID, because one feels that the PET protecting one’s privacy is easy to use. Factor 3 is characterized by two highly loading items referring to ‘helplessness’ (H 1, 2). Factor 2 is characterized by the items classified as ‘information control’ as well as one question treating contingency (CON 1). Looking into the question text for the contingency item we can interpret the loading as respondents’ perception of their PET

as a means or channel information and then determine further steps. Consequently, we feel comfortable to regard factor 2 as a control dimension that measures to what extent one perceives control as a consequence of being informed.

Tables 3 and 4 show that the cumulative variance explained by these three factors is above 78% for both PET conditions. And, it is important to note that the three factors are almost not correlated. This means that they are measuring independent dimensions of perceived control.

The final step was to investigate the internal consistency of the three scales thus identified. For this purpose we calculated each item set's Cronbach  $\alpha$ . The threshold of .8 was passed by the ease-of-use construct as well as the information control construct. The two items on helplessness displayed a rather weak Cronbach  $\alpha$  of around .6. Potentially, these questions would need to be retested in future research and be complemented with other items to form a better scale.

**Table 3.** Control scales group 1 (Password), reliability statistics

Control scales	Item	Cron $\alpha$	Culm. Variance explained	Corr (r)	Corr (r)	Corr (r)
Ease-of-use of the PET	EUP 1	.881	38,33%	.243	.110	-.214
	EUP 2					
	EUP 3					
Information Control	CON 1	.837	64,30%			
	IC 1					
	IC 2					
Helplessness	H 1	.650	78,63%			
	H2					

**Table 4.** Control scales group 2 (Agent), reliability statistics

Control scales	Item # item rank	Cron $\alpha$	Culm. Variance explained	Corr (r)	Corr (r)	Corr (r)
Ease-of-use of the PET	EUP 1	.915	34,70%	.092	.118	.050
	EUP 2					
	EUP 3					
Information Control	CON 1	.836	61,91%			
	IC 1					
	IC 2					
Helplessness	H 1	.579	78,99%			
	H2					

## 5 Applying Control Scales to UC PETs

Typically, scales identified on the basis of one sample should not be applied to the same sample for the report of actual findings. Still, in order to add practical meaning

to the control scales discussed in this article, we want to apply them here to demonstrate their usefulness.

As outlined above, we want to use the scales to measure peoples' perceived control over UC technology once they have a PET to protect their privacy. Thus, we want to measure how people perceive exercising privacy with the help of a PET. As described in section 4.2. 128 subjects answered to the control scales described above upon seeing a film on RFID deployments in retail and at home. Group 1 and 2, however, differed with respect to the PET displayed to them in the film stimulus. With this experimental set-up it became possible to test whether people perceive different levels of control depending on the type of PET used. Recall that in the user model, people would get immediate control over when to access the intelligent infrastructure. Only upon reception of a personal password the intelligent infrastructure would be able to read out peoples' RFID chips. On the other hand, the agent model proposed a PET residing on a mobile network and operating automatically on the basis of privacy preferences specified in advance of transactions. Here, control would be delegated to an agent. The hypothesis we had upon designing the experiment was that participating subjects would perceive more control in the user model and less control in the agent model. Thus, producing an argument for more research efforts in UC technology designs putting control physically into peoples' hands.

Peoples' perception of control on the basis of having one of the two PETs at their disposition is displayed in table 5. It turns out that – against expectations- perceived control is similar for both PET technologies. More specifically, people report to feel helpless (out of control) no matter what PET is at hand. This is despite the fact that they consider both PETs to be rather easy to use. The degree to which they feel informed to actively control the environment is judged on as medium. The mean control judgements indicate that the password scheme may be slightly easier to use, but this difference is statistically non significant. The conclusion that can be drawn from these results is that *no* PET presented to participants in the current study seems to induce in people a perception of control. The proposal of either PET solution must be questioned seen that people do not feel in control with any one of the two and may therefore question the ability to effectively protect their privacy with them.

## 6 Conclusion

The current article documents the development of three scales that are able to measure peoples' perception of control over being accessed when moving in UC environments and having a PET to protect their privacy. Control is measured with a view to whether people feel informed and are able to use the PET. Furthermore, loss of control is considered by the degree of helplessness perceived by users. When researchers of UC conceive technologies today that impacts peoples' privacy, they may want to test whether the environments they envision induce a positive feeling of control. The scales presented here, may serve this purpose. Especially the two factors relating to ease-of-use and information control could be used as design guidelines for UC developers.

Applying the control scales to two PET scenarios envisioned by UC scholars show that both of them do not win peoples' trust. More precisely, they do not induce a feeling of control and thus privacy. Since they are broadly the most prevalent PET options for RFID technology thought of today, this may cause designers of RFID technology to potentially rethink the marketability of the privacy processes they currently envision.

**Table 5.** Control scales and mean answers applied to two UC PETs

<b>Control Scale</b>	<b>Questions</b> (1=fully agree, 5=do not agree at all)	<b>mean</b> <b>(user</b> <b>model)</b>	<b>mean</b> <b>(agent</b> <b>model)</b>
Ease-of-use of the PET	To learn to use <the PET> would be easy for me.	1.65	2.02
	It would be easy for me to learn skillful use of <the PET>. <sup>1)</sup>	1.92	2.15
	I find <the PET> easy to use.	2.16	2.44
Information Control	Using <the PET> I would always know when and by whom I have been read out.	2.96	2.85
	Through <the PET> I would always be informed of whether and how the intelligent environment recognizes me.	2.72	2.51
	Thanks to <the PET> I could determine myself whether or not I'll interact with the intelligent environment.	2.49	2.44
Helplessness	I could imagine that if the intelligent environment set out to scan me, it would be able to do so despite <the PET>.	1.57	1.53
	Eventually <the PET> will not be able to effectively protect me from being read by the intelligent environment.>	1.92	1.78

## 6 References

- [1] F. Mattern, "The Vision and Technical Foundations of Ubiquitous Computing," *Upgrade*, vol. 2, pp. 2-6, 2001.
- [2] M. Weiser, "The Computer for the 21st Century," in *Scientific American*, vol. 265, 1991, pp. 94-104.
- [3] I. Altman, *The environment and social behavior: Privacy, personal space, territory, crowding*. Monterey, California: Brooks/Cole, 1975.
- [4] F. Schoeman, *Philosophical Dimensions of Privacy*. Cambridge, UK: Cambridge University Press, 1984.
- [5] S. Margulis, "Privacy as a Social Issue and Behavioral Concept," *Journal of Social Issues*, vol. 59, pp. 243-261, 2003.

- [6] J. H. Smith, S. Milberg, J., and S. Burke, J., "Information Privacy: Measuring Individuals' Concerns About Organizational Practices," *MIS Quarterly*, vol. 20, pp. 167-196, 1996.
- [7] M. Boyle, "A Shared Vocabulary for Privacy," presented at Fifth International Conference on Ubiquitous Computing, Seattle, Washington, 2003.
- [8] A. Juels, R. Rivest, and M. Szydlo, "The Blocker Tag: Selective Blocking of RFID Tags for Consumer Privacy," presented at 10th Annual ACM CCS 2003, 2003.
- [9] M. Langheinrich, "A Privacy Awareness System for Ubiquitous Computing Environments," presented at 4th International Conference on Ubiquitous Computing, UbiComp2002, Göteborg, Sweden, 2003.
- [10] D. Engels, R. Rivest, S. Sarma, and S. Weis, "Security and Privacy Aspects of Low-Cost Radio Frequency Identification Systems," presented at First International Conference on Security in Pervasive Computing, SPC 2003, Boppard, USA, 2003.
- [11] S. Engberg, M. Harning, and C. Damsgaard Jensen, "Zero-knowledge Device Authentication: Privacy & Security Enhanced RFID preserving Business Value and Consumer Convenience," presented at Second Annual Conference on Privacy, Security and Trust, New Brunswick, Canada, 2004.
- [12] S. Spiekermann and O. Berthold, "Maintaining privacy in RFID enabled environments - Proposal for a disable-model," in *Privacy, Security and Trust within the Context of Pervasive Computing*, vol. 780, *The Kluwer International Series in Engineering and Computer Science*, P. Robinson, H. Vogt, and W. Wagealla, Eds. Vienna, Austria: Springer Verlag, 2004.
- [13] V. Bellotti and A. Sellen, "Design for Privacy in Ubiquitous Computing Environments," presented at 3rd European Conference on Computer Supported Cooperative Work ECSCW'93, Milan, Italy, 1993.
- [14] S. Spiekermann and H. Ziekow, "RFID: a 7-point plan to ensure privacy," presented at 13th European Conference on Information Systems (ECIS), Regensburg, 2005.
- [15] E. Skinner, "A Guide to Constructs of Control," *Journal of Personality and Social Psychology*, vol. 71, pp. 549-570, 1996.
- [16] M. E. P. Seligman, *Helplessness: On Depression, development, and death*. San Francisco: Freeman, 1975.
- [17] L. Y. Abramson, M. E. P. Seligman, and J. D. Teasdale, "Learned helplessness in humans," *Journal of Abnormal Psychology*, vol. 87, pp. 49-74, 1978.
- [18] E. Langer, *The Psychology of Control*. Beverly Hills: Sage Publications, 1983.
- [19] H. Heckhausen, *Motivation and Action*. Berlin: Springer Verlag, 1991.
- [20] S. Fiske and S. Taylor, *Social cognition*. New York: McGraw-Hill, 1991.
- [21] J. Rodin, "Control by any other name: Definitions, concepts and processes," in *Self-directedness: Cause and effects throughout the life course*, J. Rodin, C. Schooler, and K. W. Schaie, Eds. Hillsdale: Erlbaum, 1990, pp. 1-15.
- [22] A. Bandura, "Human agency in social cognitive theory," *American Psychologist*, vol. 44, pp. 1175-1184, 1989.
- [23] F. Davis, "Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology," *MIS Quarterly*, vol. 13, pp. 319-348, 1989.
- [24] V. Venkatesh, "Determinants of Perceived Ease of Use: Integrating Control, Intrinsic Motivation, and Emotion into the Technology Acceptance Model," *Information Systems Research*, vol. 11, pp. 342-365, 2000.
- [25] Y. Inoue, "RFID Privacy Using User-controllable Uniqueness," presented at RFID Privacy Workshop, Massachusetts Institute of Technology, Cambridge, MA, USA, 2004.
- [26] C. Floerkemeier, R. Schneider, and M. Langheinrich, "Scanning with a Purpose - Supporting the Fair Information Principles in RFID Protocols," presented at 2nd International Symposium on Ubiquitous Computing Systems, Tokyo, Japan, 2004.

- [27] G. Churchill and D. Iacobucci, *Marketing Research: Methodological Foundations*: South-Western College Pub, 2001.
- [28] B. Rohrmann, "Empirische Studien zur Entwicklung von Antwortskalen für die sozialwissenschaftliche Forschung," *Zeitschrift für Sozialpsychologie*, vol. 9, pp. 222-245, 1978.

# Privacy & Personalization: Preliminary Results of an Empirical Study of Disclosure Behavior

Evelien Perik<sup>1</sup>, Boris de Ruyter<sup>2</sup>, Panos Markopoulos<sup>1</sup>

<sup>1</sup> Eindhoven University of Technology, Department of Industrial Design,  
P.O.Box 513, 5600 MB Eindhoven, The Netherlands  
{e.m.perik@tue.nl, p.markopoulos}@tue.nl  
<sup>2</sup> Philips Research Eindhoven, Group Media Interaction,  
Prof.Holstlaan 4 (WY 2.01), 5656 AA Eindhoven, The Netherlands  
boris.de.ruyter@philips.com

**Abstract.** This paper describes empirical research into the privacy preferences and behaviors of individuals regarding personalization in music recommender systems. These phenomena concern music recommendations based on two different types of user information: preferences for music genres and personality traits. Our results indicate similar disclosure behavior by users for both types of personal information. This contradicts attitudes of users as reported in post-experiment questionnaires and interviews. Factors found to influence disclosure behavior are: information about the purpose of the disclosure and recipients of the information, the degree of confidentiality of the information involved, and the benefits people expect to gain from disclosing personal information.

## 1 Introduction

Personalized services rely for their operation on appropriate and sufficient information about the user. This could, for example, include information about the identity of the user, usage of the service, preferences and dislikes of the user, or even exact dates and times the service has been used by the user [5]. The acquisition and storage of such information could be regarded as intrusive. Information may be collected by explicit (and conscious) input of the user, or information may be collected implicitly, without the users' explicit intervention. Especially information that is collected implicitly can lead to privacy concerns, as it may be done without sufficient awareness or control by the individuals concerned [6]. Little is known about the actual perception of users, their preferences and needs when they are exposed to such situations. This may be due to the fact that privacy is a difficult concept to study. For instance, it seems that the privacy concerns people report in surveys do not match their actual behavior [9].

This paper describes preliminary results of empirical research into the factors influencing the trade-off between the perceived benefits of personalization and the privacy 'costs' experienced by individuals. The study described in this paper is an extension of the pilot study as described in [8]. At the time of writing this paper data collection has finished, but the analysis phase is not yet complete. In our experiment participants used a purpose-built music recommendation service over the Internet and were con-

fronted with actual privacy dilemmas. We investigated people's disclosure behavior relating to two types of personal information: preferences for different music genres and information about their personality. In the study users were offered the choice to disclose information about themselves in return for improved recommendations. Their choices and stated opinions regarding this system provide us insight into the relative sensitivity of both types of information.

## 2 Method

A within-subject design was chosen for the experiment to compare the disclosure choices made by each participant in different situations. We studied the users' reactions in providing two different types of information in exchange for the music recommendation: preferences for music genres and personality traits. Research by Rentfrow and Gosling [6] has found a relation between music preferences and personality traits. This enabled us to recommend music based on the user's personality traits. For each of these information types, three different uses of the information were possible: personal use of the information only, collaborative filtering (i.e. sharing personal information with a system that matches users based on this information), and directly showing information to other users so that these users can recommend music to each other. This resulted in six consecutive situations that participants experienced. At four different points in this study each participant was offered a choice of the level of privacy he/she would choose for their personal information. The choices made by participants were logged by the system. Through post-experiment questionnaires and interviews we tried to understand the reasons behind their choices.

**Participants:** The participants were recruited by e-mail announcements via secretaries and on bulletin boards within Philips and Eindhoven University of Technology. Announcements were sent to over 1000 employees and students of the two organizations. In total 48 participants, of which 8 female and 40 male, completed the study. The ages of participants ranged from 17 to 49 with an average age of 23 years old.

**Apparatus and Materials:** The participants had access to the music recommender service through a web browser. This service offers personalized playlists of songs. Using streaming technology, these songs were made available for playing on their personal computer. The experimental recommender service is built on a database of nearly 6000 songs and covers 12 different music genres. The recommender was made especially for the purpose of the experiment to ensure the different privacy preferences were enforced at the appropriate stage of the experiment.

**Measures:** Several types of data were collected during the study. Participants' preferences for music genres were obtained by the Short Test of Music Preferences, STOMP [6]. An inventory of personality traits was made by the Ten Item Personality Inventory, TIPI [4]. A quality rating for each recommended playlist was obtained. The titles and artists of all recommended songs were logged, as well as the time and date the service was used by participants. And finally, in the on-line questionnaire a combination of open and multiple-choice questions were answered. Among others these

questions aimed to establish the privacy attitudes of individuals, their attitudes towards taking risks, and to get explanations for their disclosure behavior.

**Procedure:** People who were interested in participating were sent an e-mail with instructions. They were not told in advance that the research was about privacy concerns; they were only told in general terms that we were investigating their experience of using a personalized music recommender. The order in which participants experienced the two recommender systems was counterbalanced. Half of the participants started using the recommender system based on preferences for music genres, followed by the recommender system based on personality traits. The other half of the participants used the recommender systems in reversed order.

On the first day of using any of the two recommender systems participants were offered no choice regarding the disclosure of their personal information. The disclosure level was set by default to allow only personal use of their information. During the second and third day participants were offered a choice. Participants could choose between three levels of permission: no permission (no disclosure), restricted permission (disclosure of the information in an anonymous way) or full permission (disclosure of the information in identifiable way). On the second day this concerned using personal information for collaborative filtering, and on the third day it concerned showing it directly to other users. If participants would choose the ‘no disclosure’ option, then the music recommender would perform as on the previous day. If participants chose to disclose their information the recommendations would improve. In order to deliver benefits to users in a predictable manner the same recommender technique was used in all three cases of the two recommender systems. However, the percentage of songs recommended according to the user profile would change in the various conditions. Participants were required to rate the quality of every playlist. This was done in order to verify that the quality of the playlist has been considered by the participant, and that it has indeed improved as intended.

In the questionnaire participants were asked to comment on their choices during the experiment and to indicate what level of permission they would choose if they could choose again. Some of the participants were contacted for an interview appointment to discuss their choices in more detail. In total 21 interviews were held.

### 3 Results

#### 3.1 Disclosure of information to the system

Participants were asked to provide two types of profile information during the experiment: music preferences and personality traits. In both cases all participants provided this profile information. In the questionnaire participants were asked to explain how they felt about having to provide these types of profile information (see Table 1). Overall, it seems that participants had fewer reservations to disclose music preferences

compared to personality. Some participants questioned the usefulness of having to provide personality traits; this was not the case for music preferences.

**Table 1.** Amount of participants using specific explanation to express how they felt about providing music preferences or personality traits

<b>Explanation</b>	<b>Preferences</b>	<b>Personality traits</b>
OK	18	13
Expected / Logical / Useful	12	
No problem	6	8
No problem - if anonymous	2	2
Preferred	1	
Difficult	4	9
Question usefulness		9
Surprised / Unexpected		4
Interesting		2
(No explanation)	5	1
<b>Total</b>	<b>48</b>	<b>48</b>

**Table 2.** The amount of participants choosing a particular level of disclosure per situation during the experiment

	<b>Compare MPrefs</b>			<b>Show MPrefs</b>			<b>Compare Traits</b>			<b>Show Traits</b>		
	No	Anon	Iden	No	Anon	Iden	No	Anon	Iden	No	Anon	Iden
<b>M-T</b>	0	11	13	0	13	11	0	12	12	0	11	13
<b>T-M</b>	1	13	10	0	11	13	0	11	13	0	12	12
<b>Total</b>	1	24	23	0	24	24	0	23	25	0	23	25

Table 2 displays the amount of participants that chose a particular level of disclosure per situation. The table shows the choices made by participants per experimental sequence (M-T: music preferences-traits; T-M: traits-music preferences), as well as for the total group of participants.

As shown in table 2, the “no disclosure” option is not adopted by anyone. Furthermore, two main groups can be distinguished based on disclosure choices. One group chose anonymous disclosure in all choice situations; the other group chose disclosure including identity information in all choice situations. These groups are about equal in size. It turns out that most participants stick to their initial choice throughout the 4 choice situations. Out of the 48 participants 41 participants stuck to their initial choice (either anonymous disclosure or disclosure including identity information). Only 7 participants chose different levels of disclosure across the 4 choice situations. Out of these 7 participants who chose different levels of disclosure in the 4 choice situations, 3 participants indicated that they chose a different level of disclosure compared to the other choice situations, because they wanted to check the effect of a different choice.

Afterwards, in the questionnaire participants were asked what level of disclosure they would choose in each of the choice situations, if they could choose again (see Table 3). In total 29 participants would not change their opinion about the level of

disclosure Table 4 shows that more participants chose the ‘no disclosure’ option after the experiment. Also more people chose a lower level of disclosure for the choice situations concerning personality traits than during the experiment. Seven participants would now choose a lower level of disclosure for the choice situations involving personality traits. 3 participants would now choose a lower level of disclosure for the choice situations where information is directly shown to other users. Also 3 participants would now choose a lower level of permission in all four choice situations as compared to the choices made during the experiment. Other participants chose otherwise different levels compared to before.

**Table 3.** The amount of participants choosing a particular level of disclosure per situation after the experiment

	Compare MPrefs			Show Mprefs			Compare Traits			Show Traits		
	No	Anon	Iden	No	Anon	Iden	No	Anon	Iden	No	Anon	Iden
<b>M-T</b>	1	12	11	1	10	13	4	12	8	5	13	6
<b>T-M</b>	1	14	9	0	13	11	1	12	11	3	12	9
<b>Total</b>	2	26	20	1	23	24	5	24	19	8	25	15

### 3.2 Comparison of disclosure behavior & privacy attitudes

If the observed disclosure behavior of participants is compared to their answers in the on-line questionnaire and the information provided in the interviews a mismatch is found between thoughts or feelings and actual behavior. Information about personality traits is considered more personal than preferences for music genres. 43% of the participants felt it was worrying if other people would get access to their information about personality traits, and 50% of the participants worried about this in the case of a music content provider. Yet all of these participants gave permission to the system for the comparison of their profile information to that of others, or for directly showing it to others. Many participants comment on the perceived difference between personality traits and preferences, like this participant: *“I think (...) that those personality traits are more confidential, than (...) the one with those preferences. (...). I think it tells more about yourself.”*

This discrepancy between disclosure behavior and privacy attitudes is in accordance with the experiment of Spiekermann et al [9] and is also described by Acquisti and Grossklags [2]. The experiment by Spiekermann is conducted in the context of e-commerce. The study described in this paper, found similar results in the context of personalized applications. The discrepancy is all the most striking considering that in our experiment self-report was obtained post-hoc and was referring directly to the choices offered rather than to general attitudes and opinions.

### 3.3 Factors influencing disclosure behavior

In the on-line questionnaire and interviews the reasons behind participants' disclosure behavior were discussed. Participants mentioned, for example, the influence of the available information, especially the amount and clarity of the information was mentioned. Two participants indicated in the on-line questionnaire not to be sure of the consequences of choosing disclosure including identity, and as a result chose anonymous disclosure instead.

Another factor was the purpose or usage of the information. Participants expressed worries about not knowing how their information will be used by the system. *"I think it is important to know that if I allow someone to just go ahead, what is actually going to happen. Who knows where you will all end up, where you will be associated with...and where they are going to use [your information] for."* Some participants also questioned the relevance of providing a name along with the profile information. This relates to the type of e-commerce users identified by Spiekermann et al [9], whose privacy concerns focus on the revelation of identity related information such as name, address or e-mail.

Participants expressed worries about who gets access to their personal information: For example, because they don't know what other people may do with the information: *"If other people can link my name to certain personal information en I don't know what they can do with it or want to do with it, then I am careful with it."*

Participants do consider how sensitive information is to them before deciding to disclose this information or not. *"Because personality traits are something, (...) I think that is fairly personal. (...) The fact that I think I am extravert, or introvert, that is something completely different than when you tell someone you like 'dance' for example. (...) That is a really different kind of information that you release"*. Some participants indicate that other people cannot derive much from knowing music preferences, whereas for personality traits people may judge you before they actually get to know you. The three factors mentioned above are in line with the model of Adams and Sasse [3] who identify *Information Receiver*, *Information Usage* and *Information Sensitivity* as three critical factors for shaping privacy behavior.

Some participants also consider what benefits they will gain from disclosing the information. *"My general opinion in that respect is that as long as something does not get extremely personal, they are allowed to know everything about me, as long as I gain from it myself"*. Based on the analysis so far, it seems that different groups of people may be distinguished based on their perception of a certain situation. Some participants seem to focus mainly on potential benefits of disclosing information whereas others seem to focus mainly on the potential risk or cost of disclosing information. This finding needs to be further analyzed, based on the qualitative data collected and compared to the chosen levels of disclosure.

Another influencing factor may have been the "scientific nature" of the setup. From the on-line questionnaire and the interview data it appears that participants felt quite safe disclosing personal information in the context of this experiment, even though they were actually allowing the system to show their personal information to other users.

## 4 Discussion

We have reported on preliminary results of a study into the privacy preferences and behaviors relating to personalized music recommender systems. All participants disclosed both music preference and personality information to the system and most felt relaxed about it. In the case of personality traits at least some participants were questioning the quality of the recommendations in advance.

Participants chose equal levels of disclosure across the 4 choice situations during the experiment: Possibly, the difference between the privacy sensitivity of the 4 different choice situations is negligible. However, the interview and questionnaire data do not support this interpretation. Further analysis of this data is needed. Another explanation may be that the participant did notice the difference between the choice situations, but were lead by their tendency to cooperate with the research as much as possible (resulting in disclosure for the use of their information throughout all choice situations). It may be that the mere thought of participating in a research changes one's concerns about privacy. By the choosing this specific experimental set up, in which participants actually used a music recommender service, and were even asked to show their information directly to other people, we tried to prevent this from happening. If actual, this phenomenon should cast some doubt to experimental studies on the topic of privacy and should suggest the need for triangulation with field survey data regarding actual disclosure behavior.

After the experiment participants were asked to indicate the desired level of disclosure in each of the 4 choice situations again. This time, more people choose a lower level of disclosure in the choice situations concerning personality traits after the experiment than during. Indicating towards a higher sensitivity of personality traits compared to music preferences.

The discrepancy between disclosure behavior and privacy attitudes is in accordance with other sources [9], [2]. It shows again that for a thorough understanding of people's attitudes towards privacy, extensive field studies are required (based on actual behavior). For the design of personalized service, it is very important to take this discrepancy between privacy concerns and disclosure behavior into account. If people do use a personalized service, this does not necessarily imply that they feel comfortable using it.

In the analysis so far, it seems that different groups of people may be distinguished based on their perception of a certain situation, either focusing on the potential benefits of disclosing information or focusing on the potential risk or cost of disclosing information. This will be further analyzed using the qualitative data available. It will be interesting to see how this segmentation relates to the one found by Ackerman et al [1] and its refinement by Spiekermann et al [9].

## 5 Conclusion

This paper described preliminary results only. Further analysis is under way, especially on the qualitative data that has been collected. The questionnaire data shows

that participants had fewer problems providing music preferences compared to personality traits. Based on questionnaire and interview data the impression arises that personality traits are considered more sensitive information.

Broadly, two groups of people can be distinguished based on participants' disclosure behavior: One group choosing to disclose anonymously and the other choosing disclosure including identity information. Further investigation is needed to see in what ways these groups can be identified. Throughout the study participants persisted with their initial choice regarding the level of disclosure. A possible interpretation is that the initial level of trust users feel towards these systems is crucial to the success of such systems.

The found discrepancy between disclosure behavior and attitudes of users is in accordance with studies in other domains such as e-commerce. With this study it is shown to hold for the domain of personalization as well. In the questionnaire and interviews participants mentioned various factors influencing their disclosure behavior, such as purpose or usage of the information, the recipients of the information, the sensitivity of the information involved and the expected benefits in return for disclosure.

## References

1. Ackerman, M., Cranor, L.F., Reagle, J.: Privacy in E-Commerce: Examining User Scenarios and Privacy Preferences. In: Proceedings of the ACM Conference on Electronic Commerce (EC'99), 3-5 November 1999, Denver, Colorado, 1-8.
2. Acquisti, A., Grossklags, J. Losses, gains, and hyperbolic discounting: An experimental approach to information security attitudes and behaviors. In 2nd Annual Workshop of on Economics and Information Security, 2003.
3. Adams, A., Sasse, M. A.: Privacy in multimedia communications: protecting users not just data. Proceedings of IMH HCI'01. (2001) 49-64.
4. Gosling, S. D., Rentfrow, P. J., Swann, W. B., Jr.: A very brief measure of the Big Five personality domains. *Journal of Research in Personality*. 37 (2003) 504-528.
5. Kobsa, A.: Pseudonymous yet Personalized Interaction with Websites that Utilize Network-wide User Modeling Services. In: 2003 HCIC Winter Workshop, Winter Park, CO, 2003.
6. Kobsa, A., Schreck, J.: Privacy through Pseudonymity in User-Adaptive Systems. In: *ACM Transactions on Internet Technology*, vol. 3 (2), 2003, pp. 149-183.
7. Perik, E.M., Ruyter, B. de, Markopoulos, P., Eggen, J.H.: The Sensitivities of User Profile Information in Music Recommender Systems. In: PST 2004: Proceedings of the Second Conference on Privacy Security and Trust, 13-15 October 2004, Fredericton, NB, Canada, 137-141.
8. Rentfrow, P. J., Gosling, S. D.: The do re mi's of everyday life: The structure and personality correlates of music preferences. *Journal of Personality and Social Psychology*. 84 (2003) 1236-1256.
9. Spiekermann, S., Grossklags, J., Berendt, B.: E-privacy in 2nd generation E-commerce: privacy preferences versus actual behavior. Proceedings of the 3rd ACM conference on Electronic Commerce. ACM Press. (2001) 38-47.

# Informed Consent to Address Trust, Control, and Privacy Concerns in User Profiling

Thea van der Geest, Willem Pieterse, and Peter de Vries

University of Twente, Faculty of Behavioural Sciences, Department of Communication Studies. P.O. Box, 217, 7500 AE Enschede, The Netherlands  
t.m.vandergeest@utwente.nl

**Abstract.** More and more, services and products are being personalised or tailored, based on user-related data stored in so called user profiles or user models. Although user profiling offers great benefits for both organisations and users, there are several psychological factors hindering the potential success of user profiling. The most important factors are trust, control and privacy concerns. This paper presents informed consent as a means to address the hurdles trust, control, and privacy concerns pose to user profiling

## 1 Introduction

On May 24 2005, personalisation was the prime topic in the newspaper headlines and the radio and TV news in The Netherlands. On that day, the Dutch Minister responsible for government reform informed the Parliament about a study of the administrative hurdles that local and national administrations create for the average citizen. The study showed that the collective Dutch citizens spend 112 million hours to meet administrative and bureaucratic demands, filling out forms and making regulations work for them. The number one annoyance of citizens is that they have to fill out forms with personal data that they have provided over and over again. The Minister announced his objective of reducing the administrative burden by 25 %, particularly by investing in electronic, personalised communication, services and transactions. In the next few years, the Dutch citizens will increasingly be presented with forms that are pre-filled with all personal data available, to be accessed through a citizen's personal portal. Various government departments, such as the Tax and Welfare agencies, are already working on the realisation of personalised forms, transactions and portals. Their experience shows that there is more at stake than just technical and organisational issues. How can organisations like government agencies re-use personal data in a way that is acceptable for the average citizen? This paper relates acceptance of the use of personal data in electronic communication and services to the underlying personal psychological factors of trust, control and privacy concerns.

The rise of ICT and the Internet since the 90's of the past century has led to new possibilities for the purchase or acquisition of services or products. People no longer have to visit shops or counters to get information, communicate or perform transactions. But the new technologies have more possible benefits. Personalisation

(also indicated as customisation, or tailoring) is one of those benefits. Organisations can collect data about their clients and use it intelligently for the planning and adaptation of messages, information or actions with or for the individual. In that case, the organisations use the data about current user characteristics or behaviour to adapt information and communication to the targeted individual and to predict future behaviour. Well known commercial examples of online personalisation are portals like My Yahoo (yahoo.com) or recommender systems like the online bookseller Amazon (amazon.com) has created.

Re-use of data collected or provided on earlier occasions can strengthen the relationship between user and organisation and increase the effectiveness and efficiency of communication both for the user and the organisation. A good user-experience during the contact will lead to (more) satisfaction about the application used, e.g. e-commerce or e-services, and more importantly, to a (more) positive image of the organisation behind the application [1].

In order to make ‘intelligent’ use of user-related information, that is, to personalise products and services an organisation needs to build a profile or user model of its customers or citizens. We define a user profile as follows:

*A user profile is a (structured) data record, containing user-related information including identifiers, characteristics, abilities, needs and interests, preferences, traits and previous behaviour in contexts that are relevant to predicting and influencing future behaviour [38].*

Some categories of user-related information concern stable, unalterable ‘properties’ of the user, such as name, age and gender. Other categories relate to properties that can easily alter over time (e.g. developing new preferences or abilities) and context (e.g. having a need for information during international travel, but not during national travel).

A number of social and psychological factors, however, reduce the acceptability of user profiling. A majority of users expresses privacy concerns about the use of personal data on the Internet, as will be discussed further on in this paper. This leads, for example, to websites like [www.bugmenot.com](http://www.bugmenot.com) where you can obtain a login name and password to various websites (like [nytimes.com](http://nytimes.com)) without having to register. The sense of control of about one’s own user profile is another important factor. Organisations sometimes collect and distribute data about individuals without the users knowing and wanting this. As Alpert et al. [3] show, users want to be in control. A third major factor influencing acceptance is trust. In order for user profiling to be successful, users have to trust user profiling.

Many organisations have tried to deal with trust, control and privacy issues. For example, websites try to take away privacy concerns by means of a privacy statement or seal. Other organisations offer users the possibility to control their own user profile, in line with the EU directive on the protection of personal data [16]. This paper discusses the concept of informed consent, mentioned in the directive, as a strategy to address trust, control, and privacy concerns in user profiling.

In the remainder of this paper, trust, control and privacy will be discussed more thoroughly, and then informed consent will be presented and discussed in detail.

## 2 Trust

The first of the psychological aspects influencing the acceptance of user profiling is Trust. A number of theorists have proposed trust to be a mechanism that enables people to deal with situations of uncertainty or risk. Luhmann [26], for instance, argued that trust effectively limits the number of possible behavioural outcomes associated with dealing with other people to only a relatively small number of expectations. Limiting the investigation of all possible outcomes of an interaction to only a few, may result in more careful investigation of the realistic option, which may reduce both uncertainty and risk of the actor. In a similar vein, Anthony Giddens [21] used the term trust for situations where knowledge about the other party, i.e., the trustee or referent, is absent.

In light of the above it is not surprising that trust is generally accepted as a prerequisite for good personalisation practice [6]. Users are not likely to reveal confidential information about themselves and may be suspicious of data harvesting practices if they fear that this information could be misused in some way, and that they, consequently, put themselves at risk by doing so. Research [23] demonstrated that lack of trust was the major reason for people not to engage in online shopping. In addition, Warkentin, Gefen, Pavlou, and Rose [40], found that trust in the organisation using the technology and trust in governmental policies are important determinants for the adoption of e-services. They state that trust is a crucial enabler affecting purchase intentions, inquiry intentions and the intention to share personal information. The latter intention, of course, is especially relevant in user profiling. Briggs et al. [6] point to the fact that trust and personalisation have a reciprocal relationship. Trust is not only a prerequisite for good personalisation, good personalisation also generates trust.

Trust, however, is not a unitary concept. It has been studied in various disciplines, ranging from economics and political sciences to personality research and social psychology, each of which may treat the concept differently with regard to whether trust is seen as a dependent, independent or interaction variable, whether it is static or dynamic, or whether it is studied on the institutional, group or individual level (for an overview see [5], [15], [31]). The next paragraphs discuss various forms of trust.

The concept of *general trust*, or generalised interpersonal trust, for instance, relates to the trust people have in most other people, or in strangers, and is treated as a stable characteristic of both individuals and groups [15]. As such, general trust can be seen as a necessary prerequisite for other forms of trust to develop; without a general sense of trust, a user would not be willing to enter interactions of any kind.

Contrary to general trust, *social trust* is based on social relations and shared values. The actors at which this type of trust is directed are more concrete than with general trust; specifically, they are persons or organisations that are perceived to share the trustor's values [33]. Social trust, a focus of attention in risk management research, involves little or no interaction, and is often a 'one-shot' affair [15]. Value similarity may be inferred after shooting only a quick glance at the trustee; simple cues, such as skin colour or gender may be enough for the trustor to infer that if the trustee looks similar, he or she may also hold similar values. If user profiling is aimed at establishing social trust, the profile should contain information about the relevant values that the profiled person holds about social issues, persons and organisations.

*Interpersonal trust* is established and maintained in and through interaction and communication. It is a kind of trust much studied in social psychology where it is treated as an expectation of the other's behaviour that is specific to the interaction [10]. This expectation is argued by some to be based on perceptions of the other's competence and honesty [29] or goodwill [41]. If a user profile contained the information on the basis of which interpersonal trust can be predicted, it should be fed with information about the interactions occurring between the organisation and the user. This means that the user profile needs to be updated continuously.

Different labels for and distinctions between types of trust are found in the literature of the different fields. However, most are analogous to the typology described above. Zucker [43], for instance, used the term *characteristic trust* to denote trust based on social relations, comparable with Earle et al.'s [15] concept of social trust. In addition, Rotter [30] distinguished between *dispositional* and *relational trust*, the former relating to others in general, the latter based on interaction with a particular other. *Propensity to trust*, proposed by Mayer, Davis and Schoorman [28] as a stable characteristic affecting the likelihood that someone will trust, may be thought of as a general willingness to trust others, and as such, it bears a strong resemblance to general trust.

Online interaction with an organisation involves both the organisation itself, as well as a system which enables this interaction. Obtaining tax refunds online, for instance, involves the tax agency as the organisation that enables and controls online interactions, as well as several interfaces that enable clients to submit information about their income and deductible expenses electronically. *Organisational trust* and *system trust* are, therefore, of particular importance to the implementation and acceptance of user profiling. The former is a type of trust that partly overlaps the categories of social and interpersonal trust: it has an organisation or group as its referent, as does social trust, and at the same time is based on interactions, as is typical of interpersonal trust (e.g., see Zaheer, McEvily and Perrone [42]). The latter, system trust can be seen as a special case of interpersonal trust. Like interpersonal trust it refers to expectations about behaviour of a specific other, rather than a group of others or strangers. In the case of system trust, however, the referent is not a human partner, but rather an object, i.e. the system with which a user is in interaction.

In sum, acceptance of user profiling is influenced by the user's trust propensity in general, trust in the organisation he or she is dealing with, and trust in the systems the organisation uses to interact and communicate with the user, including the user profiling system.

Trust is related to many other issues that appear to be critical for user profiling. Firstly, trust is influenced by the *sense of control* about the user profile [4]. When end users feel that they themselves or a trusted third party representing them controls the user profile and its applications, they will trust user profiling more than when they feel that the organisations in control are not primarily focusing on the users' interests, or, put differently, do not share the user's values.

Trust is also influenced by *privacy concerns*. Concern about the privacy aspects of personal information shared on the Internet is correlated with increasing levels of Internet experience [20]: the more experienced internet users are more worried about privacy issues. There is considerable resistance among many Internet users to engage in business-to-consumer transactions over the Web, primarily due to concerns about

privacy and the trustworthiness of the Internet [2], [39]. Findings of Chellappa & Sin [8] also stress the relationship between trust and privacy. In an empirical study, they found that both trust and privacy factors correlate significantly with the likelihood of using personalised services. Also they found that privacy and trust were correlated. Factors building trust (like familiarity and past experiences) led to lower privacy concerns.

### 3 Control

Alpert et al. [3] studied user attitudes regarding the personalisation of content in e-commerce websites. In their study, the users expressed their strong desire to have full and explicit control of personal data and interaction. They want to be able to view and edit (update and maintain) their personal information at any time.

A study by Roy Morgan Research [32] shows that 59% of the 1524 Australian respondents state that their trust in the Internet increases when they feel they have control over their personal information. The study also showed that:

- 91% of the respondents want to be asked for explicit permission before companies use their information for marketing purposes;
- 89% of the respondents want to know which persons and which organisations have access to their personal information;
- 92% of the respondents want to know how their personal information is used.

User control obviously is a critical condition for user acceptance of profiling and personalisation. However, the study cited does not answer the question whether the users themselves should host the user profile themselves, nor whether trusted third parties can resolve the users' anxiety about control issues.

Byford [7] perceives personal information as a property or asset of the individual ('Byford's property view'). The user is the owner of his or her personal information. In Byford's property view, individuals see privacy as the extent to which they control their own information in all types of Internet exchanges. The property aspect of the exchange manifests itself in the users' willingness to trade personal information for valued services such as free e-mail or special discounts from merchants.

A user profiling system that is not supported by a good system for user control of personal information is bound to lead to acceptance problems. However, creating the interaction and the related user interface that allow users to control the information in their profiles is a complicated problem, especially if the control goes beyond a very coarse level of granularity [11]. Although users have indicated they want to be in control of their personal data, very little users make use of possibilities websites offer to control personal information. A number of ecommerce web sites give users access to their profiles. However, it is not clear whether users are aware of this facility [11, p.69]. Reports of operators of personalisation systems have indicated that users rarely take actions to proactively customise their online information [27].

## 4 Privacy concerns

Violation of privacy is one of the most important concerns of Internet users. As much as 70 – 84 % of all participants in various surveys indicated that privacy concerns made them resist providing personal data. They are especially aware of privacy issues concerning personal data, such as name, addresses and income. Also, 24-34 % of people in the surveys indicated to have provided false or fictitious information, when asked to register [12], [18], because of concerns about privacy violation. In commercial contacts (on-line shopping) those privacy concerns play an even more important role than in other systems for tailoring information or communication. As much of 91 % of respondents indicated that they were concerned about businesses sharing user data for purposes other than the original purpose for collecting the data [37]. Although many Internet users are not well-informed about the means of collecting usage data (web surfing behaviour data), such as spyware and cookies, almost everybody (91%) indicates to feel uncomfortable about being tracked across websites [22].

All these figures indicate that privacy and personal data protection are of the utmost importance to almost all Internet users. However, this does not mean that they understand the implications of their concerns and act upon it. Only 10% of respondents in a survey had their browsers installed in such a way that it rejected cookies [18]. In a study of Spiekermann et al. [34] even users with self-reported strong privacy concerns readily disclosed personal and sensitive information on a web site. Although people express concern about privacy, they easily give up on privacy because of convenience, discounts and other incentives, or a lack of understanding of the consequences. Obviously there is a difference between concerns and attitudes at one hand and actual secure behaviour at the other hand.

Yet, the privacy concerns of users imply that organisations should approach the process of user profiling with extreme caution. Effective user profiling depends on the correctness of information and on the willingness of user to provide data to the organisation. Technical solutions such as good privilege regulations, and regulatory solutions like required privacy policies, could help to secure privacy and thus to reduce privacy concerns. But we assume that they will not work without accompanying measures to address the users' attitudes. Creating trust, giving users control, and requesting informed consent are essential conditions for solving the privacy issue. The organisation, as the initiator of collecting user data and user profiling, should take the initiative to protect and secure the users' privacy. Common practice of organisations is to add a privacy statement to websites and consider the problem solved. As Kobsa & Teltzrow [24] show, privacy statements are hardly read, let alone comprehended by the visitors of websites. Their research shows that the design of privacy statements (like the user interface) might help in alleviating users' privacy concerns.

Loeb [25] distinguishes three types of privacy concerns: regarding protection of the user profiles and queries, regarding protection of the person's web usage history and regarding protection of the actual information if the delivery takes place over public networks.

Wang et al. [39] distinguish four types of privacy threats:

- improper acquisition of information (e.g. uninvited tracking of the users' web usage);
- improper use of information (e.g. distribution of data to third parties);
- privacy invasion (e.g. spamming a mailbox with uninvited direct mailings);
- improper storage and control of personal information (e.g. no opting-out, no means to remove incorrect or unwanted information)

It is still unclear which of the privacy threats and concerns mentioned are (most) influential for acceptance of user profiling. But an overview of studies regarding privacy and personalisation on the Internet does show that users have significant concerns over the use of personal information for personalisation purposes on the Internet [36]. CyberDialogue [13] found that 82% of all Internet users say that a website's privacy policy is a critical factor in their decision to purchase online. Even more salient is that 84% of the respondents have refused to provide information at a website because they were not sure how that information would be used.

The fact that there is a concern, however, does not necessarily imply that users don't provide any information. The lack of trust in privacy policies moved a large majority of users to give false or fictitious information over the Internet, and thus protect their privacy [12], [37]. According to research conducted by the Winterberry Group, this development is increasingly becoming a problem for the collection of user-related information [14]. Two aspects of data quality seem to interfere here: whether the personal data collected are an accurate, adequate and reliable representation of the user, and whether they are used in an acceptable, appropriate and allowable way.

## **5 An integral solution to address control, trust and privacy concerns: Informed consent**

Trust, control and privacy are strongly related concepts. Paying attention solely to establishing a trustworthy relationship between users and the supplier of personalised services and products is fruitless when no attention is paid to privacy and control issues. It might well be possible that a user does trust the organisation offering personalisation, but feels his privacy is being threatened when supplying personal information and therefore does not use the personalised e-services of that organisation. For example; Chellappa and Sin [8] found that "while vendors can do little to positively influence consumers' concern for privacy directly, our analysis sheds light on the possibility for them to indirectly affecting consumers' privacy concerns through trust building". On the other hand, enabling users to exert control over their own information, may increase trust and, thus, reduce privacy concerns. Informed consent, a mechanism that enables users to exert control, is a requirement under the EU Data Protection Directive of 1995 (95/46) and the subsequent directive 2002/58 on privacy and electronic communications.

The 1995 and 2002 directives describe consent of a user as "any freely given specific and informed indication of his wishes by which the data subject signifies his agreement to personal data relating to him being processed." The consent seems to be defined in a negative way: it offers protection from intrusion. The 2002 directive even

suggests that consent can be dealt with sufficiently by ticking off a checkbox on a web site. This conception of informed consent does not create possibilities for building and managing trust, or for users who see their personal information as a commodity they want to use for identity management with interested organisations.

In the health care sector informed consent on the use and application of personal data is defined more extensively, compared to the definitions of the EU directives. Patients have the legal and ethical right to be informed about what will happen to their body, and make informed decisions about the intervention or treatment before it is started.

*Informed consent is the process by which a fully informed user participates in decisions about his or her personal data. It originates from the legal and ethical right the user has to direct what happens to his or her information, and from the ethical duty of organisations using personal data to involve the user in the control, use and maintenance of these data.<sup>1</sup>*

Sreenivasan [35] argued that informed consent in medicine consists of two parts: a duty to obtain the voluntary agreement of patients or trial participants before treatment or enrolment, and a duty to disclose adequate information to the patient or participant before seeking this agreement.

According to Friedman, Millet and Felten [19], informed consent in Web privacy policies comprises *disclosure, comprehension, voluntariness, competence, and agreement*.

*Disclosure* refers to providing accurate information about the benefits and harms that might reasonably be expected from the action under consideration. What is disclosed should address the important values, needs and interests of the individual. *Comprehension* refers to the individual's accurate interpretation of what is being disclosed. This component raises the question: What criteria must be satisfied in order to say that something has been adequately comprehended? For example: does a user understand the privacy statement? Why (not)? *Voluntariness* means that an individual only should participate voluntarily, there may be no control about an individual's actions and the action may not be coerced. *Competence* refers to possessing the mental, emotional and physical capabilities needed to be capable of giving informed consent. Children, for example, might not be mentally and emotionally capable to judge whether or not to provide personal information on websites. Finally, *agreement* refers to a reasonably clear opportunity to accept or decline to participate [19]. This not only implies the opportunity to choose whether or not to participate at all, but also to the opportunity to choose to stop or continue the participation at any time. This means, for user profiling, that the individual should have the full control at all time.

In both Sreenivasan's [35] and Friedman et. al.'s [19] work, control and information are considered central to informed consent. Accurate information about the potential benefits and harm, provided by competent sources, should aim at comprehension of that information, and, thus, pave the way for voluntary agreement with the suggested treatment of the patient (or disagreement). The same should apply to organisations operating outside the medical world, such as providers of e-services.

---

<sup>1</sup> See: <http://eduser.v.hscer.washington.edu/bioethics/topics/consent.html>.

These organisations have much to gain from implementing informed consent in their online interactions with customers. When users are informed whether personal information will be collected, and, if so, how it will be used, they can assess the chances of their information-sharing leading to unpleasant consequences of any kind, which greatly reduces their uncertainty.

As is illustrated by the examples mentioned in the above, the effect of informed consent is twofold. First, users are informed about the procedures that will be employed by the organisation they are dealing with. This, in fact, reduces the uncertainty that users may experience when engaging in online interactions, and, consequently reduces the need for high levels of trust (cf. [26], [9]). Furthermore, informed consent enables users to exert control and retain ultimate responsibility over what they feel is sensitive information. This may well cause users to think less negatively about information gathering, and, more importantly, the intentions of the organisation: by implementing informed consent, organisations may communicate that they have their users' best interests at heart, which would greatly increase user trust in the organisation (cf. *social trust*, [15], [33]).

When informed consent is perceived and realized as a process that involves users in the control of their personal data, it offers promising perspectives for an integral strategy to deal with trust, and privacy concerns, thus increasing acceptance.

Based on the experiences with informed consent in the field of medicine, we propose that the following elements should be addressed in an informed consent procedure regarding user profiling.

1. The nature of the personal data collected for the sake of user profiling.
2. The organisation's objectives with user profiling and its prospective effects for the user. This includes the sharing of data with other organisations, and their respective objectives for user profiling (cross-domain user profiling).
3. The alternatives when no data are collected, or when no user profiling is applied. Also, the alternatives when particular types of user-related information are rejected, or when particular applications of user profiling are refused.
4. Relevant risks, benefits and uncertainties related to user profiling, for the various alternatives.
5. Assessment of the user's understanding of the information.
6. Explicitly stated acceptance or declining by the user, for all or particular types of user-related information, and for all or particular applications of user profiling.

The consent must be voluntary, and the user must have the competence to understand the information and its consequences, or the right to decide on the use of one's own personal information is void. Therefore, special attention must be paid to those groups in society that do not have easy access to ICT. Both the procedure and the information on user profiling should be explained in layperson terms. The user's understanding and acceptance must be assessed along the way, not only at initial adoption of user profiling.

Informed consent is a critical condition from the perspective of the individual user, but it might not always be in the interest of organisations to inform the public about the collection and use of user-related information. According to Business Week<sup>2</sup> 88% of users want sites to garner their consent when personal information is collected.

---

<sup>2</sup> See: [http://www.businessweek.com/2000/00\\_12/b3673010.htm](http://www.businessweek.com/2000/00_12/b3673010.htm).

According to a report from the Federal Trade Commission, 59% of websites that collect personal identifying information neither inform Internet users that they are not collecting such information nor seek the user's consent [17]. This strongly conflicts with the public's interest and is a violation of European privacy and personal information protection laws.

Informed consent requires efforts from organisations; they have to start a dialogue with their users about e.g. the control of the user profile. Organisations have to inform their users about their privacy status and the consequences of engaging in user profiling. Benefits, however, are numerous; users are well informed and are able to make proper decisions, raising levels of trust, assuring privacy and dealing with the control of the user profile. Little empirical data is available that deals with informed consent and user profiling. It is necessary to research the dimensions of informed consent. What factors determine informed consent? Do people understand consent? When do we call someone "informed"? Do people oversee the consequences of their consent? Both qualitative and quantitative research methods might be used to explore the dimensions of informed consent and the many ways in which it can be presented to the public and their effect on control and privacy concerns. Such studies can help us to assess the impact of control, trust and privacy issues on user profiling.

## References

1. Accenture: Leadership in Customer Service: New Expectations, New Experiences. Accenture, (2005)
2. Aldridge, A., Whithe, M., Forcht, K.: Security considerations of doing business via the Internet: cautions to be considered. *Internet Research*, 7(1) (1997) 9-15
3. Alpert, S.R., Karat, J., Karat, C.-M., Brodie, C., Vergo, J.G.: User attitudes regarding a User-Adaptive eCommerce Web Site. *User Modelling and User-Adapted Interaction*, 13(4) (2003) 373-396
4. Araujo, I., Araujo, I.: Developing trust in Internet commerce. In: 2003 conference of the Centre for Advanced Studies on Collaborative research. Toronto, Canada, (2003)
5. Bhattacharjee, A., Devinney, T.M., Pillutla, M.M.: A formal model of trust based on outcomes. *Academy of Management Review*, 23(1998) 459-472
6. Briggs, P., Simpson, B., De Angeli, A.: Personalisation and Trust: A reciprocal Relationship? In: Karat, C.-M., Blom, J.O. and Karat, J. (eds.): *Designing Personalized user experiences in eCommerce*. (2004)
7. Byford, K.S.: Privacy in Cyberspace: constructing a model of privacy for the electronic communications environment. *Rutgers Computer and Technology Law Journal* (24) (1998) 1-74
8. Chellappa, R.K., Sin, R.: Personalization versus Privacy: An Empirical Examination of the Online Consumers' Dilemma. *Information technology and management*, 6(2-3) (2005)
9. Coleman, J.S.: *Foundations of social theory*. Harvard University Press, Cambridge (1990)

10. Corritore, C.L., Kracher, B., Wiedenbeck, S.: Online trust; concepts, evolving themes, a model. *International Journal of Human-Computer studies*, 58(2003) 737-758
11. Cranor, L.F.: I Didn't buy it for myself: Privacy and ecommerce personalization. In: Karat, C.-M., Blom, J.O. and Karat, J. (eds.): *Designing Personalized user experiences in eCommerce*. Kluwer Academic Publishers, Dordrecht (2004)
12. Culnan, M.J., Milne, G.R.: The Culnan-Milne survey on consumers & online privacy notices: Summary of Responses. In: *Interagency Public Workshop: Get Noticed: Effective Financial Privacy Notices*. Washington DC, (2001)
13. CyberDialogue: Online consumer personalization survey. The personalization consortium, Wakefield (2001)
14. Direct Marketing: Anonymous Web Browsing Threatens Profiling Practices of E-marketers. (2001)
15. Earle, T.C., Siegrist, M., Gutscher, H.: Trust and confidence: A dual-mode model of cooperation. Western Washington University, WA, USA (2002)
16. European Union: Charter of Fundamental Rights of the European Union (Article 8(1)). Nice (2000)
17. Federal Trade Commission: Privacy online: Fair information practices in the electronic marketplace. (2000)
18. Fox, S., Raine, L., Horrigan, J., Lenhart, J., Spooner, T., Carter, C.: Trust & Privacy Online: Why Americans want to rewrite the rules. The Pew Internet & American Life Project, Washington DC (2000)
19. Friedman, B., Millet, L., Felten, E.: Informed consent online: A conceptual model and design principles. UWCSE Technical Report, 00-12-2(2000)
20. George, J.F.: Influences on the Intent to make Internet purchases. *Internet Research*, 12(2) (2002) 165-180
21. Giddens, A.: *The consequences of modernity*. Stanford University Press, Stanford, CA (1990)
22. Harris Interactive: A Survey of consumer privacy attitudes and behaviors. Harris, Rockester, NY (2000)
23. Hoffman, D.L., Novak, T.P., Peralta, M.: Building consumer trust online. *Communications of the ACM*, 42(4) (1999) 80-85
24. Kobsa, A., Teltzrow, M.: Contextualized Communication of Privacy Practices and Personalization Benefits: Impacts on Users' Data sharing and Purchase Behavior. In: Martin, D. and Serjantov, A. (eds.): *Privacy Enabling Technologies*, Springer Verlag Lecture Notes in Computer Science. (2004)
25. Loeb, S.: Architecting personalized delivery of multimedia information. *Communications of the ACM*, 35(12) (1992) 39-47
26. Luhmann, N.: *Trust and Power: Two works by Niklas Luhmann*. John Wiley & Sons, Chichester (1979)
27. Manber, U., Patel, A., Robinson, J.: Experience with personalization on Yahoo! *Communications of the ACM*, 43(8) (2000) 35-39
28. Mayer, R.C., Davis, J.H., Schoorman, F.D.: An integrative model of organizational trust. *Academy of Management Review*, 20(1995) 709-734
29. Renn, O., Levine, D.: Credibility and trust in risk communication. In: Kasperson, R.E. and Stallen, P.J.M. (eds.): *Communicating risks to the public*. Kluwer, Dordrecht (1991) 175-218

30. Rotter, J.B.: Interpersonal Trust, trustworthiness, and gullibility. *American Psychologist*, 35(1980) 1-7
31. Rousseau, D.M., Sitkin, S.B., Burt, R.S., Camerer, C.: Not so different after all: A cross discipline view of trust. *Academy of Management Review*, 23(1998) 393-404
32. Roy Morgan Research: Privacy and the community. (2001)
33. Siegrist, M., Cvetkovich, G.T., Gutscher, H.: Shared Values, social trust, and the perception of geographic cancer clusters. *Risk Analysis*, 21(2001) 1047-1053
34. Spiekermann, S., Grossklags, J., Berendt, B.: E-privacy in 2nd generation E-commerce: Privacy preferences versus actual behavior. In: *ACM Electronic Commerce 2001 conference*. (2001) 38-47
35. Sreenivasan, G.: Does informed consent to research require comprehension. *The Lancet*, 362(December 13) (2003) 2016-2018
36. Teltzrow, M., Kobsa, A.: Impacts of User Privacy preferences on personalized systems: a comparative study. In: Karat, C.-M., Blom, J.O. and Karat, J. (eds.): *Designing personalized user experiences for eCommerce*. Kluwer Academic Publishers, Dordrecht (2004)
37. UMR: Privacy Concerns Loom Large. Study Conducted for the Privacy Commissioner of New Zealand. (2001)
38. van der Geest, T.M., van Dijk, J.A.G.M., Pieterse, W.J. (eds.): *Alter Ego: State of the art on user profiling. An overview of the most relevant organisational and behavioural aspects regarding User Profiling*. Telematica Instituut, Enschede (2005)
39. Wang, H., Lee, M.K.O., Wang, C.: Consumer Privacy concerns about Internet marketing. *Communications of the ACM*, 41(3) (1998) 63-70
40. Warkentin, M., Gefen, D., Pavlou, P.A., Rose, G.M.: Encouraging citizen adoption of e-Government by building trust. *Electronic Markets*, 12(3) (2002) 157-162
41. Yamagishi, T., Yamagishi, M.: Trust and commitment in the United States and Japan. *Motivation and Emotion*, 18(1994) 130-166
42. Zaheer, A., McEvily, B., Perrone, V.: Does trust matter? Exploring the effects of interorganizational and interpersonal trust on performance. *Organizational Science*, 9(1998) 141-159
43. Zucker, L.G.: Production of trust: Institutional sources of economic structure 1840-1920. In: Staw, B.M. and Cummings, L.L. (eds.): *Research in organizational behavior*. JAI Press, Greenwich, C.T. (1986) 53-111

# A Software Product Line Approach for Handling Privacy Constraints in Web Personalization<sup>1</sup>

Yang Wang and Alfred Kobsa

Donald Bren School of Information and Computer Sciences  
University of California, Irvine, U.S.A.  
{yangwang, kobsa}@ics.uci.edu

**Abstract.** Web personalization has demonstrated to be advantageous for both online customers and vendors. However, its benefits are severely counteracted by privacy concerns. Personalized systems need to take these into account, as well as privacy laws and industry self-regulations that may be in effect. When these constraints are present, they not only affect the personal data that can be collected, but also the methods that can be used to process the data. The present research aims at maximizing the personalization benefits, while at the same time satisfying the currently prevailing privacy constraints. Since such privacy constraints can change over time, we seek a systematic and flexible mechanism that can cater to this dynamics. We looked at several existing approaches and found that they fail to present a practical and efficient solution. Inspired by the ability of software product lines to support software variability, we propose a user modeling architecture based thereon that supports architectural level configuration management to dynamically select personalization methods that satisfy current privacy constraints. A pilot experiment is being carried out with the support of an existing user modeling server and a software architecture based development environment.

## 1 Introduction

Personalization technologies have been successfully introduced on the World Wide Web where they are mostly used for customer relationship management [1]. A number of studies show that personalization has provided benefits for both online customers and vendors [2, 3].

However, personalization benefits are offset by privacy concerns [4-7]. Since personalized websites collect personal data, they are also subject to privacy laws and regulations if the respective individuals are in principle identifiable. A review of nearly 30 international privacy laws [8] shows that if privacy laws apply to a personalized website, they often not only affect the data that are collected by the website and the way in which data is transferred (e.g., to which party), but also the methods that may be used for processing them (and consequently the components that embed such methods). For instance, the German Teleservices Data Protection Act [9] mandates personal data to be erased immediately after each session except for very

---

<sup>1</sup>This research has been supported through NSF grant IIS 0308277. We would like to thank André van der Hoek and Eric Dashofy for their helpful comments.

limited purposes. This provision could affect the use of machine learning methods where the learning takes place over several sessions.

From a personalization point of view, we ask the research question: how can personalized web-systems maximize the personalization benefits, while at the same time being compliant with the privacy constraints that are currently in effect (such as privacy laws, industry and company regulations, and the privacy preferences of the current user)? The remainder of this article is organized in the following way: we will discuss several existing approaches in Section 2, our proposed software product line approach in Section 3, our pilot experiment in Section 4, our example in Section 5, and finally present conclusions in Section 6.

## **2 Existing Approaches**

### **2.1 Anonymous Personalization**

Basically, this approach allows users to remain anonymous with regard to the personalized system and the whole network infrastructure, whilst enabling the system to still recognize the same user in different sessions so that it can cater to her individually [10]. At first sight, this seems to be the panacea because in most cases privacy laws do not apply any more when the interaction is anonymous. However, anonymity is currently difficult and/or tedious to preserve when payments, physical goods and non-electronic services are being exchanged, it harbors the risk of misuse, and it hinders vendors from cross-channel marketing (e.g. sending a products catalog to a web customer by mail). Moreover, users may still have additional privacy preferences (e.g., they do not want profiling even when it is only done pseudonymously), to which the personalized system needs to adjust.

### **2.2 Largest Permissible Dominator**

Ideally, this approach means that only those personalization methods that meet all privacy laws and regulations are used. The Disney website for instance meets both the U.S. Children's Online Privacy Protection Act (COPPA) as well as the European Union Directive [11]. This solution is likely to run into problems if more than a very few jurisdictions are involved, since the largest permissible denominator may then become very small.

### **2.3 Different Country/Region Versions**

In this approach, personalized systems have different country versions, with personalization methods only that are admissible in the respective country. If countries have similar privacy laws, separate versions can be built for these countries

combined, using the above-described largest permissible denominator approach. For example, IBM's German-language pages meet the privacy laws of Germany, Austria and Switzerland [12], while IBM's U.S. site meets the legal constraints in U.S. This approach is also likely to be infeasible as soon as the number of countries/regions, and hence the number of different versions of the personalized system, increases.

### 3 Our approach

User modeling systems are widely used for supporting user-adaptive applications. In industry, most such systems use a client-server architecture. A User Modeling Server (UMS) stores and represents user characteristics and behavior, integrates external user-related information, applies user modeling methods to derive additional assumptions about the user, and allows multiple external user/client adaptive applications to retrieve user information from the server in parallel [13]. Since privacy constraints directly affect UMSSs, we suggest addressing them in the UMS design.

#### 3.1 A Dynamic Privacy-Enabling User Modeling Architecture

For many personalization goals, more than one method can often be used that differ in their data and privacy requirements and their anticipated accuracy and reliability. For example, a personalized website could use incremental machine learning (that discards all raw data after the end of a session) to provide personalization to web visitors from Germany<sup>2</sup>, while it can use possibly better one-time machine learning with the data stored across several sessions to provide personalization to web visitors from the U.S. who are not subject to this constraint. We propose a software architecture that encapsulates different personalization methods in individual components and, at any point during runtime, ascertains that only those components can be operational that are in compliance with the currently prevailing privacy constraints. Moreover, the architecture can also dynamically select the component with the optimal anticipated personalization effects among those that are currently permissible. To implement this design, we choose the Software Product Line (SPL) approach from software architecture research.

#### 3.2 Software Product Line Architecture (PLA)

Software Product Lines have been successfully introduced in industrial software development for improving productivity, software quality and time-to-market [14]. A product line architecture represents the architectural structure for a set of related products by defining *core elements* that are present in all product architectures, and *variation points* where differences might occur among specific product architectures.

---

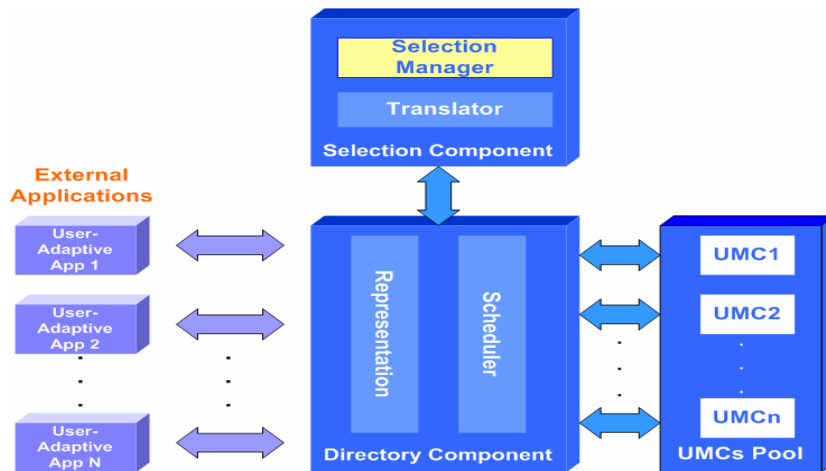
<sup>2</sup> This is not yet a complete solution though since the German Teleservices Data Protection Act also mandates that profiling requires the use of pseudonymous or the consent of the user.

Each variation point is guarded with a Boolean expression. Given a set of desired properties or bindings (expressed in name-value pairs), a particular product architecture can be selected out of a product line architecture by resolving the Boolean guards of each variation point.

Treating software as a product line is a new approach to support software variability from design-time to invocation-time to run-time [15]. We conceive our user modeling server as a product line architecture, where each personalization component (embedding a specific personalization method) forms a variation point in the architecture. Components that are in compliance with the currently prevailing privacy constraints will be flagged, and only those will be able to operate. If required, the architecture can additionally select the component with the highest personalization benefits based on a designer-specified preference order.

## 4 Pilot Experiment

We are conducting a feasibility study based on an existing user modeling server and an architecture-based software development environment. Figure 1 shows our privacy-enabling user modeling architecture implemented as a product line architecture.



**Fig. 1.** A Dynamic Privacy-Enabling User Modeling Architecture

### 4.1 ArchStudio 3.0 – An architecture-base development environment

Our user modeling architecture is implemented in Archstudio 3.0 [16], which supports architecture-level configuration management (such as versioning, diff and merge operations) and deployment of product line architectures. We express the UMS in

xADL 2.0 [17], the underlying XML-based architectural description language for ArchStudio 3.0.

## 4.2 External User-Adaptive Applications

On the left side of Figure 1, we see several external user-adaptive applications that query user information from the UMS in order to provide personalization services, and supply new information about the user. They communicate with the UMS using standard LDAP operations such as add, search, bind and unbind.

## 4.3 Directory Component

We used an existing LDAP-based<sup>3</sup> common user data repository [18] as our test bed for managing user models. The repository is represented as *Directory Component* in Figure 1 and is composed of two sub-systems: *Representation* and *Scheduler*. The Representation sub-system is in charge of managing directory content (i.e., mainly user-related information). The Scheduler sub-system is responsible for the communication between the Directory Component and various *User Modeling Components* (UMC).

## 4.4 User Modeling Components (UMC)

On the right side of Figure 1, we see a set of user modeling components. Each of these components embodies a user modeling method (e.g., collaborative filtering, domain-based inferences). A UMC can subscribe to certain types of internal events of Directory Component by maintaining event subscriptions in the Service Model hosted by the Representation sub-system. After the launch of the UMS, the Scheduler loads event subscriptions from the Service Model. Subsequently, the Scheduler periodically checks the Service Model for new entries and, if necessary, updates its internal subscription tables accordingly. Henceforth, the Scheduler acts as an event broker that supervises LDAP events within the Directory Component and communicates them together with associated data to UMCs.

In our product line architecture, UMCs are treated as variant components guarded by Boolean expressions which express privacy constraints<sup>4</sup> pertaining to the personalization methods incorporated in these components.

## 4.5 Selection Component

*The Selection Component* which is shown in the top part of Figure 1 subscribes to the

---

<sup>3</sup> Our approach can use any form of user data repository, e.g., a database, as long as personalization methods can plug into the repository.

<sup>4</sup> While ultimately these constraints should be expressed in privacy constraint specification languages (such as APPEL [19] or EPAL [20]) or semantic web technologies [21], we currently use a set of variables only.

LDAP events of the Directory Component. The Selection Component comprises two sub-components: *Selection Manager* and *Translator*. The main use of the Selector Manager is to carry out the dynamic user modeling component selection described in Section 4.6. The Selection Manager is implemented in a component-based and message-based architectural style called C2 [22]. The Translator maps the LDAP events to C2 messages. Figure 2 shows the internal architecture of the Selection Manager. The xArchADT stores the architectural description of the PLA. The Selector performs the selection of personalization methods, and the Manager orchestrates the whole selection and instantiation process.

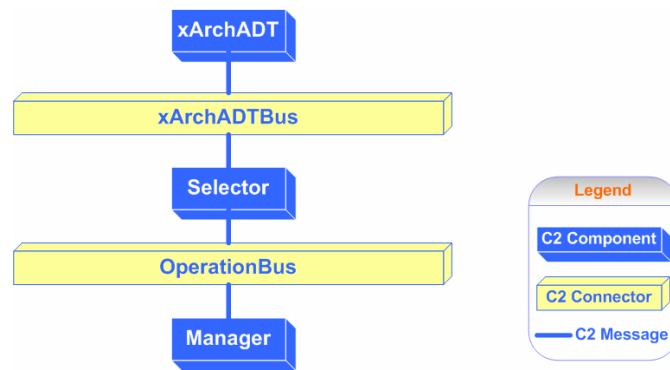


Fig. 2. Internal Structure of the Selection Manager

#### 4.6 Dynamic Selection Mechanism

The Manager monitors the start and end of user sessions via external applications' LDAP bind and unbind operations onto the UMS. When the Manager detects the start of a user session, it initiates a *Privacy Context Detection* process that will collect all the active privacy constraints and then generate corresponding variable bindings for the privacy constraints of all UMCs. A similar process will be carried out whenever during a user session the Manager learns about new or changed privacy requirements (which for all practical purposes will stem from user preferences since privacy laws and regulations are unlikely to change during a user session).

The bindings are fed into the Selector that will carry out a *PLA selection process*. First, the Boolean guards of all UMCs are evaluated based on their variable bindings, to determine whether or not these UMCs may be included in the user modeling architecture for the current user session. A binary Privacy Constraint Satisfaction (PCS) vector is constructed whose  $n^{\text{th}}$  element represents whether or not the  $n^{\text{th}}$  UMC may be used. The Selector checks whether a run-time system instance with such a PCS already exists. If so, the Manager will assign the user session to the existing run-time system instance that has the same PCS. If not, the Selector will perform *PLA Pruning* that automatically removes any disallowed components from the architecture, and then the Manager instantiates a new run-time system instance for the user session.

The following pseudo-code illustrates the above dynamic selection mechanism:

The Selection Manager monitors LDAP bind and unbind events of user sessions:

On bind:

*Privacy Context Detection:*

Collect the current privacy constraints (e.g., privacy laws and regulations);  
Generate variable bindings for Boolean guard expressions of UMCs;

*PLA selection*, based on variable bindings:

Evaluate Boolean guards for UMCs;  
Construct a new Privacy Constraints Satisfaction (PCS) vector V where element V[i] signifies whether UMC<sub>i</sub> may be used (1) or not (0) for the user session;

**IF** there already exists an identical PCS **THEN**

Assign the user session to the existing run-time system instance;  
*numOfSessions* ++; //add 1 to the number of sessions handled by this instance

**ELSE**

*PLA Pruning:*

Prune out UMCs whose Boolean guards are resolved to FALSE, meaning they are in conflict with the privacy constraints in effect;

Instantiate a new run-time system instance for the user session;  
*numOfSessions* =1; // assign the user session to the new run-time instance

On unbind:

*numOfSessions* - -; // decrease the number of sessions handled by this instance by 1

**IF** *numOfSessions* == 0 **THEN**

Kill the corresponding run-time system instance;

If new/changed user privacy preferences are detected:

Similar process as on bind, but simplifications are possible;

## 5 Example

UniversalFriends.com is a website run by UniversalFriends LLC in the USA, a signatory of the U.S. Network Advertisers Initiative. The goal of this website is to bridge physical distances between people and to foster universal friendship via information technology. It provides personalized services to help customers make friends worldwide. Upon registration, each user will be asked to choose a pseudonymous user ID along with a password and provide some information about themselves (e.g., their hobbies). Users will be given some space on the UniversalFriends web server to create their own homepages. Based on a user's characteristics, the system will recommend a personalized list of likely friends, and will automatically send invitations for pairwise virtual meetings.

The UniversalFriends web server relies on our privacy-enabling user modeling server to provide inferred information about users to recommend potential friends. More specifically, inferences about a user are calculated by different user modeling components from the User Modeling Components Pool as shown in Table 1.

User Modeling Component	Data used	Methods used
UMC <sub>1</sub>	<ul style="list-style-type: none"> <li>Demographic data (age, gender, profession, education level and so forth)</li> </ul>	Clustering techniques using demographic data (e.g., recommend people in the same profession cluster).
UMC <sub>2</sub>	<ul style="list-style-type: none"> <li>User-supplied data (e.g., a user indicates her levels of interests in different topics)</li> </ul>	Rule-based reasoning (e.g., if a user indicates a high interest in a specific topic, we infer that she would like to meet people with similar ratings for the topic).
UMC <sub>3</sub>	<ul style="list-style-type: none"> <li>User-supplied data</li> </ul>	Fuzzy reasoning with uncertainty (e.g., if a user indicates a high interest in a specific topic, we are 95% confident to infer that she would like to meet people with similar ratings for the topic).
UMC <sub>4</sub>	<ul style="list-style-type: none"> <li>Demographic data</li> <li>User-supplied data</li> </ul>	Rule-based reasoning
UMC <sub>5</sub>	<ul style="list-style-type: none"> <li>Demographic data</li> <li>User-supplied data</li> </ul>	Fuzzy reasoning with uncertainty
UMC <sub>6</sub>	<ul style="list-style-type: none"> <li>User-supplied data</li> <li>The UniversalFriends pages the user visited in the current session</li> </ul>	Incremental machine learning
UMC <sub>7</sub>	<ul style="list-style-type: none"> <li>User-supplied data</li> <li>The UniversalFriends pages the user visited across sessions</li> </ul>	One-time machine learning
UMC <sub>8</sub>	<ul style="list-style-type: none"> <li>Demographic data</li> <li>User-supplied data</li> <li>The UniversalFriends pages the user visited across sessions</li> </ul>	One-time machine learning, Fuzzy reasoning with uncertainty

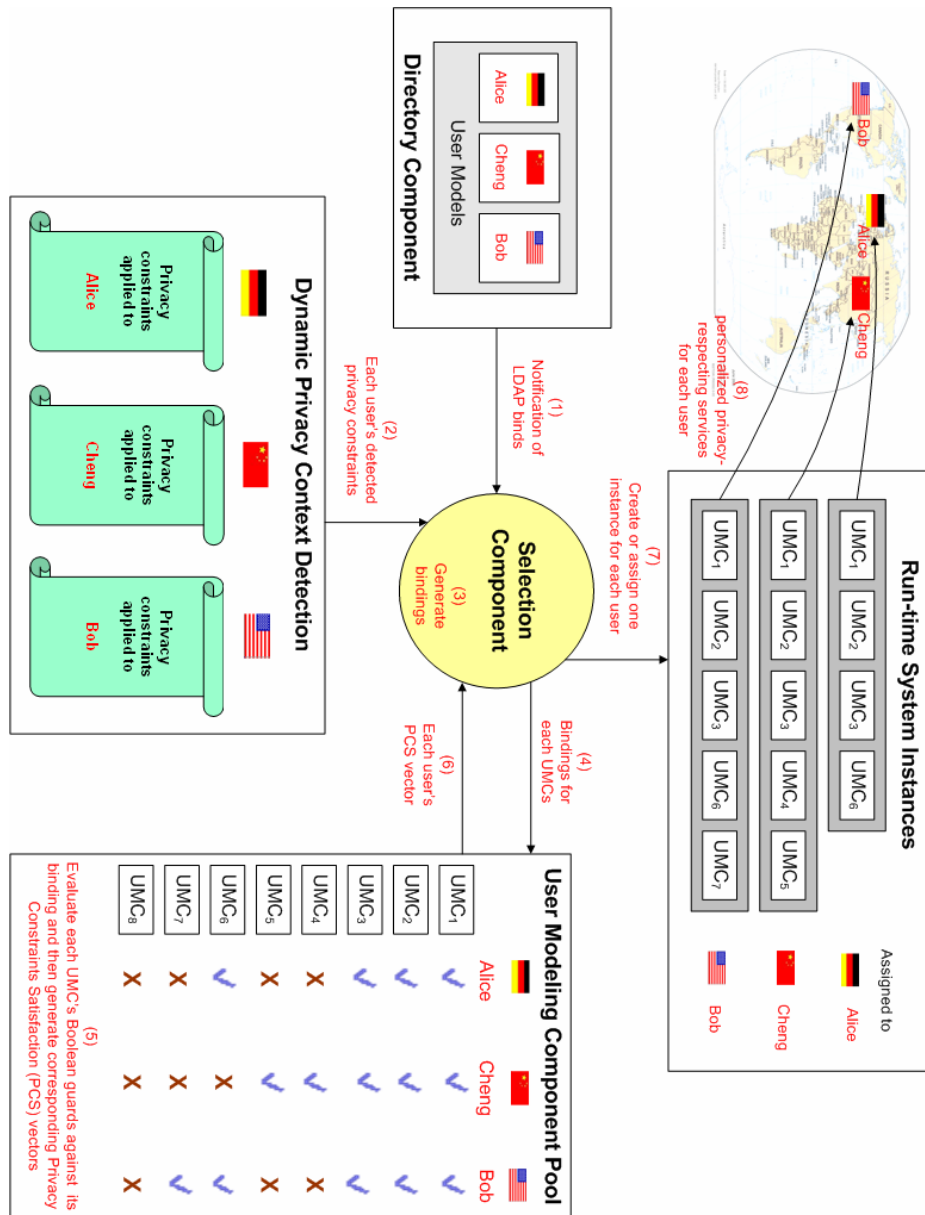
**Table 1.** Different User Modeling Components in the User Modeling Components Pool

Let us assume that we have three users: Alice, Cheng and Bob. Table 2 describes their information:

Name	Nationality / Current Location	Privacy preference
Alice	Germany	None
Cheng	China	Dislike being tracked
Bob	The United States	None

**Table 2.** Information about our hypothetical users

Figure 3 illustrates the process of how our privacy-enabling user modeling architecture caters to each individual user.



**Fig. 3** Privacy-enabling User Modeling Process

As is shown in Figure 3, the user modeling process is controlled by the Selection Component. Assume that the three users have requested the website for recommendations of potential friends. The web server will bind to the Directory Component (an LDAP server in our implementation) and then send relevant user information to the individual user models. The Selection Component will check users' privacy constraints via a privacy context detection, which in turn will generate the bindings of each user modeling component for every user (i.e., whether or not it can be used according to the user's privacy constraints), and filter out the UMCs that are not allowed to operate.

The privacy constraints that apply to each of the three individual users and their implications to the UMCs are discussed below:

For Alice, Germany's Tele-Services Data Protection Law applies:

- UMC<sub>4</sub>, UMC<sub>5</sub>, UMC<sub>8</sub> are illegal because the law prohibits combining user profiles retrievable under pseudonyms with data relating to the bearer of the pseudonym.
- UMC<sub>7</sub>, UMC<sub>8</sub> are illegal because the law mandates personal data to be erased immediately after each session except for very limited purposes.

Therefore, UMC<sub>4</sub>, UMC<sub>5</sub>, UMC<sub>7</sub>, UMC<sub>8</sub> cannot be used for Alice.

Despite no privacy law that can apply to Cheng was found, she has her own personal privacy preference as "dislike being tracked". UMC<sub>6</sub>, UMC<sub>7</sub>, UMC<sub>8</sub> cannot be used because the system cannot keep track of the pages she visits on UniversalFriends.com.

For Bob from the United States, UMC<sub>4</sub>, UMC<sub>5</sub> and UMC<sub>8</sub> cannot be used according to the NAI self-regulation [23] if he does not give consent on merging non-personally identifiable use data with personally identifiable demographic data.<sup>5</sup>

To provide privacy-enhanced personalized services to users, the Selector Component will produce a PCS vector and instantiate a new run-time system instance for each user, or use an existing instance if its PCS is the same as that of another user.

## 6 Conclusions

Our approach facilitates the construction of personalized websites operate in a privacy-aware manner (both with respect to legal and user requirements). Our software product line approach allows personalized websites to address the combinatorial complexity of privacy constraints in a systematic and flexible manner, which builds on state-of-the-art industry practice for managing software variants at runtime. We aim at exploring the feasibility of this approach using an existing user modeling server and empirically established privacy constraints.

---

<sup>5</sup> We can choose a single method from the set of permissible methods by using a partial selection preference order derived through domain analysis at design time.

## References

1. Kobsa, A., J. Koenemann and W. Pohl, *Personalized Hypermedia Presentation Techniques for Improving Online Customer Relationships*. The Knowledge Engineering Review, 2001. **16**(2): p. 111-155.
2. Hof, R., Green, H., and Himmelstein, L., *Now it's YOUR WEB*, in *Business Week Oct. 5*. 1998. p. 68-75
3. Personalization Consortium, *Personalization & Privacy Survey*. 2000, Personalization Consortium: Edgewater Place, MA.  
<http://www.personalization.org/SurveyResults.pdf>
4. FOR, *The Privacy Best Practice*. 1999, Forrester Research: Cambridge, MA
5. IBM, *IBM Multi-National Consumer Privacy Survey*. 1999, IBM.  
[http://www.ibm.com/services/files/privacy\\_survey\\_oct991.pdf](http://www.ibm.com/services/files/privacy_survey_oct991.pdf)
6. DePallo, M., *AARP National Survey on Consumer Preparedness and E-Commerce: A Survey of Computer Users Age 45 and Older*. 2000, AARP: Washington, D.C. <http://research.aarp.org/consume/ecommerce.pdf>
7. Teltzrow, M. and A. Kobsa, *Impacts of User Privacy Preferences on Personalized Systems: a Comparative Study*, in *Designing Personalized User Experiences for eCommerce*, C.-M. Karat, J. Blom, and J. Karat, Editors. 2004, Kluwer Academic Publishers: Dordrecht, Netherlands.
8. Chen, Z. and A. Kobsa, *A Collection and Systematization of International Privacy Laws, with Special Consideration of Internationally Operating Personalized Websites*. 2002. <http://www.ics.uci.edu/~kobsa/privacy>
9. DE-TS, *German Teleservices Data Protection Act*. 1997.  
[http://www.datenschutz-berlin.de/recht/de/rv/tk\\_med/iukdg\\_en.htm#a2](http://www.datenschutz-berlin.de/recht/de/rv/tk_med/iukdg_en.htm#a2)
10. Kobsa, A. and J. Schreck, *Privacy through Pseudonymity in User-Adaptive Systems*. ACM Transactions on Internet Technology, 2003. **3**(2): p. 149-183.
11. *Personal Communication, Chief Privacy Officer, Disney Corporation*. 2002
12. *Personal Communication, Chief Privacy Officer, IBM Zurich*. 2003.
13. Kobsa, A., *Generic User Modeling Systems*. User Modeling and User-Adapted Interaction, 2001. **11**(1-2): p. 49-63.
14. Bosch, J., *Design and Use of Software Architectures: Adopting and Evolving a Product-Line Approach*. 2000, New York: Addison-Wesley
15. van der Hoek, A., *Design-Time Product Line Architectures for Any-Time Variability*. Science of Computer Programming, special issue on Software Variability Management, 2004. **53**(30): p. 285-304.
16. ArchStudio, *ArchStudio 3.0*. 2005. <http://www.isr.uci.edu/projects/archstudio/>
17. Dashofy, E.M., A.v.d. Hoek, and R.N. Taylor. *A Highly-Extensible, XML-Based Architecture Description Language*. in *Proceedings of The Working IEEE/IFIP Conference on Software Architecture*. 2001. Amsterdam, Netherlands.: IEEE Computer Society. 103-112
18. Fink, J., *User Modeling Servers: Requirements, Design, and Evaluation*. 2004, IOS Press, Netherlands (Infix)
19. Cranor, L., M. Langheinrich, and M. Marchiori, *A P3P Preference Exchange Language 1.0 (APPELI.0): W3C Working Draft 15 April 2002*. 2002.  
<http://www.w3.org/TR/P3P-preferences>

20. Schunter, M. and C. Powers, *The Enterprise Privacy Authorization Language (EPAL 1.1): Reader's Guide to the Documentation*. 2003: IBM Research Laboratory. <http://www.zurich.ibm.com/security/enterprise-privacy/epal/>
21. Gandon., F.L. and N.M. Sadeh, *Semantic Web Technologies to Reconcile Privacy and Context Awareness*. *Journal of Web Semantics*, 2004. 1(3): p. 241-260. doi:10.1016/j.websem.2003.07.008
22. Taylor, R.N., et al, *A Component- and Message-Based Architectural Style for GUI Software*. *IEEE Transactions on Software Engineering*, 1996(22(6), June, 1996): p. P.390-406.
23. NAI, *The Network Advertising Initiative (NAI) Self-Regulatory Principles*. 2001, Network Advertising Initiative

# Privacy and Security in Ubiquitous Personalized Applications

Ajay Brar, Judy Kay

School of Information Technologies  
University of Sydney  
{abrar1, judy}@it.usyd.edu.au

**Abstract.** Personalization systems provide customized service based on user preferences. In ubiquitous computing environments, personalization can be achieved based on user preferences stored on mobile devices. This requires a mechanism for capturing user information and making it available to users. However, storing and exchanging potentially personal information raises user privacy concerns. Past solutions to privacy in ubiquitous computing have been ad-hoc, application specific or partially implemented. This work explores a general framework to providing privacy-aware personalization in ubiquitous computing environments. A prototype implementation of the framework has been developed and evaluated. This paper describes the approach used and its underlying concepts.

## 1. Introduction

An important current development in the area of computing is the emergence of ubiquitous computing. Ubiquitous computing (or ubicomp) envisions a computational environment integrated into the physical world, featuring a multitude of heterogeneous computing devices interacting seamlessly to provide information when and where required. This proliferation of computing into the physical world suggests new paradigms of human computing interaction inspired by constant access and increase in information and computational capabilities [1]. Personalized services form one such interaction.

Personalization in ubiquitous computing environments would depend on detecting user characteristics and preferences and providing services based on these. User preferences may be stored on the user's mobile device, e.g. a PDA, and released in exchange for personalized services. Service providers would benefit from being able to provide improved and differentiated services, such as, targeted advertising and loyalty reward programs. Users would benefit from receiving information and services customized to their preferences. For example, a user could walk into a video store and receive information about latest releases in genres of his liking along with news about special offers that match his viewing preferences. The genre preferences and the user's viewing habits would be stored in a user model on his PDA and would form the basis of any personalization offered by the service providers, in this case, the video store.

User models provide a natural construct for storing, managing and communicating user information to personalize services in ubiquitous computing environments. However, maintaining user models and releasing them to service providers, raises privacy and security concerns relating to the storage, access and processing of the information contained within the model. Users may wish to limit and control the amount of information released to the service provider. The information stored is often personal and may potentially identify the user. Protecting user identity would require providing mechanisms to allow anonymous/pseudonymous interaction with personalized services. Service providers would also need to ensure the accuracy and authenticity of the information provided by the user. Mechanisms are thus required for defining relevant subsets of the user model, allowing users to control their release, and authenticating and protecting the integrity and confidentiality of the user information released to the service provider. These topics have been extensively researched in relation to the World Wide Web [8, 10], but have received much less attention [5] in the context of ubiquitous computing.

The majority of past work on privacy has focused on providing anonymity, hiding user identity and keeping personal information secret [7]. However, this addresses a narrow aspect of privacy and does not cover scenarios where users *want* to share selective information with others [7]. Participation in the social world requires disclosure of information; users need to provide information about them in order to personalize information and services. Thus, privacy needs to be seen in terms of a negotiated controlled disclosure of information. This paper addresses a wider notion of privacy that focuses on providing users with notice of data collection, choice regarding collection and informed consent, so they can make informed decisions regarding the disclosure of their personal information.

In this paper, we propose a framework for providing personalized services/information to users, based on user preferences stored in mobile devices while minimizing risks to user privacy. The framework, called Secure Persona Exchange (SPE), provides a set of tools for capturing, representing and storing user preferences as subsets of a user model, releasing the information to a user's mobile device, and restricting and controlling the release of the information to service providers while protecting user privacy.

## 2. Related Work

This work draws on research in the areas of user modeling and privacy mechanisms for the Web and in ubiquitous computing. User modeling has an important role in ubiquitous computing and is essential for personalization of user environments through user information (collected from sensors) stored in user models [9]. The information including location information, current activity, preferences etc can be used to provide tailored services to users. Kay et al describe the architecture of distributed, ubiquitous user models [9]. The architecture relies on the concept of partial user models, called *personas*, which are accessed by services in the ubiquitous environment. Personas implemented using the scaled-down user-modeling server, PersonisLite, can be stored on user mobile devices. Users can control the content of

each persona by defining access control at different levels: evidence source, component, view and context [9]. Personas were used to represent user information in the SPE framework.

The Platform for Privacy Preferences (P3P) [13] is a specification by the World-Wide Web Consortium (W3C) and provides a framework for informed online interactions on the Web, allowing users to negotiate agreements with services that declare their privacy practices and request data. P3P provides a vocabulary and standard machine-readable format to allow websites to declare their privacy practices and describe the data they collect. User privacy preferences are described using A P3P Preference Exchange Language (APPEL 1.0), defined in a companion specification [14]. Users can use this language to express their preferences as a set of rules (called a ruleset), which can then be used by their user agent to make automated or semi-automated decisions regarding the acceptability of machine-readable policies from P3P compliant Web sites [14].

There has been past work on adapting P3P to provide privacy in ubiquitous computing environments. Langheinrich suggests a P3P-style architecture to provide notice of data collection in ubiquitous computing environments [11]. Sensors and other recording devices can use a P3P declaration format to announce via one or more well-known mechanisms, their data collection practices. User agents on user mobile devices can release contextual information based on a comparison between the privacy declaration and user preferences encoded in a machine-readable format similar to APPEL. The Privacy Awareness System (pawS) implements P3P in ubicomp environments to provide notice-choice based data collection [12]. The SPE framework includes a P3P style privacy awareness mechanism.

### **3. Framework Design**

The Secure Persona Exchange (SPE) framework is based on P3P with an underlying notice-choice privacy model. Personas are used to represent user information and provide personalization. Machine-readable policies based on the P3P vocabulary are used to provide notice of data collection and decisions are based on user preferences expressed in APPEL. The framework also contains provisions to allow access to services with varying degrees of anonymity. The framework and its architecture is described in detail in [24]. This section describes the high level requirements and features of the framework.

#### **3.1 End User Requirements**

Table 1 summarizes the end-user requirements for the system. These requirements are based on analysis of research papers discussing user privacy preferences in ubicomp systems [6, 7, 15]; analysis of privacy laws and regulations [3]; and analysis of design guidelines for privacy-aware ubicomp systems [2, 4, 11].

**Table 1.** Summary of end-user requirements

End – user requirements
<ul style="list-style-type: none"><li>• Purpose specification</li><li>• Openness</li><li>• Simple and appropriate controls</li><li>• Limited data retention</li><li>• Pseudonymous interaction</li><li>• Decentralized control</li></ul>

First, a mechanism is needed that allows users to view what benefits are offered by the personalized service and what personal information is needed to offer those benefits. They should be aware of the purpose of data collection [3]. This corresponds to the ‘Clear value proposition’ mentioned in [7]. The *PURPOSE* tag in P3P privacy policies can be used to describe the purpose of data collection and specify what benefits will be provided through personalized service.

Second, users should be aware of any data collection that takes place [3]. While this may seem to contradict the invisibility property of ubicomp systems, the requirement actually means that users should be able to view the data collected at any instant in time. A user agent may still perform actions on behalf of the user; but a log of all requests should be maintained and users should be able to configure the agent to prompt them for certain persona requests. This can be achieved using the PROMPT attribute of APPEL rules and corresponds to the “Notice/Awareness” principle mentioned by [11].

Third, users want simple control over the information disclosed and the entity to which this information is released [6, 7]. User information can be divided into personas based on the sensitivity of the information and the personalization, which can be provided based on it. For example, a user may have personal and public personas. Users may decide to disallow any access to the personal persona but allow any service to access the public persona. These access levels can be configured using APPEL rulesets.

Fourth, there have been user concerns over long-term retention of personal data [7]. Since, data is stored would be stored by service providers, limiting retention of personal data is outside the scope of the system. However, the P3P “RETENTION” tag can be used to discover the length of the period for which service providers would store user data and preferences can be specified to release information only to service providers with a limited data retention policy.

Fifth, users may prefer to interact with personalized services under assumed identities (pseudonyms), perhaps without divulging their actually identity. Pseudonymous interaction is supported: users can maintain a collection of user models (personas). Users can select one of these for use with a particular personalized service.

Finally users are concerned about systems that centralize data since sensitive data is stored on computers outside their control [6]. SPE follows a distributed architecture with user data stored on their mobile devices and thus under their control.

### **3.2 Personalization**

The SPE framework implements client-side personalization where user information is provided to service providers through subsets of the user model known as personas stored on the user's mobile device. A persona thus captures the information regarding the user and their preferences that is needed for personalizing a particular service. Client-side personalization through personas addresses user concerns relating to the storage of personal data on systems outside their control. It also addresses legal requirements present in privacy laws and regulations relating to the storage of personal information. A major challenge for client-side personalization is ensuring the authenticity and integrity of user information. This is realized in the SPE framework through an Authorizing Entity responsible for creating user personas and releasing them to users. The Authorizing Entity signs the persona and also separately provides the service provider with persona templates (describing the structure of the persona) to allow them to interpret the information provided by the user. The Authorizing Entity is thus the trusted third party in the exchange and plays a role similar to the Certification Authority in PKI. An issue with client side personalization is the availability of user information. Since the information release is controlled by the user, the availability of the information cannot be guaranteed. This is, however, a business issue and is not addressed in this work.

### **3.3 Notice Choice Privacy Model**

The SPE framework is based on the P3P notice-choice privacy model. Service providers issue requests for user information represented in one or more personas along with their privacy policy described using P3P vocabulary. The privacy policy describes the purpose of data collection, the entity collecting the data, all entities that shall have access to the data, the period for which the data will be stored and other statements required by P3P. This comprises the notice part. The privacy policy is evaluated against the user's privacy preferences and the user may configure these to prompt for action, release data or block data release based on the contents of the privacy policy. Default privacy preferences are provided for usability purposes but users can also define their own privacy preferences (based on the APPEL specification) and are thus not restricted to choosing the least intrusive of a set of preferences. This provides users with choice regarding data collection thus allowing them to make informed decisions regarding the release of his information. The P3P-like elements thus provide appropriate notice and user privacy preferences defined using APPEL provide appropriate choice.

### **3.4 Pseudonymous Interaction**

The SPE framework can be combined with a network-level anonymizer to provide pseudonymous access to personalized services. Personalization systems may not need to know the actual identity of the person involved. All they need is some way of distinguishing different sessions and relating a particular set of interactions to some identity and maintain this across multiple sessions. The SPE framework supports pseudonymity through the concept of personas. Users may interact with services using different personas containing separate information, add new personas to their mobile

device and switch between them. Multiple users can store their personas on the same device and personalize services based on these.

### 3.5 Mobile Context

The core requirement underlying the SPE framework is personalization within a mobile ubiquitous computing environment. Thus user information is stored on mobile devices (the prototype was implemented on a PDA) and communicated to service providers across a wireless network (the prototype used WiFi).

## 4. Security Requirements and Mechanism

Security in user modeling is not a goal in itself, but an auxiliary means for realizing privacy [16]. The same principle applies to ubiquitous computing; security measures are designed to protect the privacy of the data exchanged and the entities involved in the system.

Requirements for security comprise requirements for implementing the four key attributes of security, i.e., authentication, confidentiality, integrity and non-repudiation [19]. Another key attribute of security is availability. Within the context of user modeling, availability refers to the amount of user modeling functionality available to user model client (in this case, the service provider) [16]. Since the functionality is adjustable depending upon user preferences, as expressed in APPEL, availability cannot be guaranteed. Similarly, ubiquitous computing applications depend on the availability of networking infrastructure (such as WiFi) and thus, availability of a particular ubiquitous system cannot be independently guaranteed. Thus, availability is not included as a security requirement.

The security requirements discussed here apply only to securing the communication between the participants of the system. Service providers would be responsible for securing the storage of personas and templates on their systems. The system also does not provide mechanisms for securing personas and templates stored on the user's mobile device. A solution may be to store encrypted personas along with a one way keyed hash of the persona. On a PDA, however, this may introduce additional processing delays. Similarly, the authorizing entity is expected to provide its own methods for securing the data storage.

Requirements for communication security are discussed below:

- Confidentiality: personas may contain personal information and thus their content needs to be kept secret from entities other than the participants in the system. Secrecy of exchange can be achieved through SSL.
- Integrity: personas and templates need to be protected against tampering during communication. This may again be achieved by using secure message digests and communicating over SSL. Users would be responsible for ensuring the integrity and confidentiality of personas stored on their mobile devices.
- Authentication: users need to authenticate personas and templates released by the authorizing entity. Similarly, service providers need to authenticate the templates released by the authorizing entity and the personas released by the user. Additionally, users need to authenticate the service provider prior to releasing their personas. Thus there are two kinds of authentication that is required: *entity* authentication to authenticate the participant in the exchange and *data*

*authentication* to authenticate the personas and templates exchanged. Note that users do not authenticate themselves while communicating with service providers. This allows them to preserve their anonymity. Entity authentication can be achieved through X.509 certificates and communicating over SSL while data authentication can be achieved through RSA signatures together with a Certification Authority. Note that data authentication implicitly provides data integrity (for if a message has been modified, the source has changed) [19].

- Non-repudiation: refers to preventing an entity from denying previous commitments or actions [19]. Non-repudiation is not a core security requirement of the system but may be required to prevent a service provider from denying data collection.

## 5. Conclusions and Future Work

The SPE framework provides privacy-aware personalization in ubiquitous computing environments. We have presented an analysis of user requirements, requirements for personalization, privacy and security based on past research. We also described how the Secure Persona Exchange (SPE) framework addresses these requirements. Current challenges in the implementation relate to the immaturity of software support for PDAs, including the serious issue of fixed MAC addresses.

## References

1. Abowd, G. D. and E. D. Mynatt (2000). "Charting past, present and future research in ubiquitous computing." ACM Transactions on Computer-Human Interaction, Special Issue on HCI in the new Millenium. 7(1): 29-58
2. Adams, A. (2000). "Multimedia Information Changes the Whole Privacy Ball Fame". Proceedings of Computer, Freedom, and Privacy. Toronto, Canada. ACM Press.
3. Australian Privacy Act (1988). "Information Privacy Principles under the Privacy Act 1988."
4. Belotti, V. and A. Sellen (1993). "Design for Privacy in Ubiquitous Computing Environments". Proceedings of the Third European Conference on Computer Supported Cooperative Work (ESCW'93). Milan, Italy. Kluwer Academic Publishers.
5. Hitchens, M., J. Kay and B. Kummerfeld (2004). "Secure identity management for pseudo-anonymous service access". University of Sydney School of Information Technologies Technical Report TR 546. June 2004
6. Hong, J. I., G. Boriello, J. A. Landay, D. W. Mc Donald, B. N. Schilit and J. D. Tygar (2003). "Privacy and Security in the Location-enhanced World Wide Web". Proceedings of Fifth International Conference on Ubiquitous Computing: Ubicomp 2003 (Workshop on Ubicomp Communities: Privacy as Boundary Negotiation). Seattle, WA
7. Hong, J. I. and J. A. Landay (2004). "An Architecture for Privacy-Sensitive Ubiquitous Computing". Proceedings of the 2<sup>nd</sup> International Conference on Mobile Systems, Applications, and Services (MobiSYS). Boston, Massachusetts, USA.

8. Kay, J., R. J. Kummerfeld and P. Lauder (2003). "Managing private user models and shared personas". UM03 Workshop on User Modelling for Ubiquitous Computing.
9. Kobsa, A. and J. Schreck (2003). "Privacy Through Pseudonymity in User-Adaptive Systems." ACM Transactions on Internet Technology. 3(2): 149-183.
10. Langheinrich, M. (2001). "Privacy by Design – Principles of Privacy-Aware Ubiquitous Systems". Ubicomp 2001 Proceedings.
11. Langheinrich, M. (2002). "A Privacy Awareness System for Ubiquitous Computing Environments". Ubicomp 2002.
12. Cranor, L. F., M. Langehinrich, M. Marchiori, M. Presler-Marshall and J. Reagle (2002). "The platform for privacy preferences 1.0 (p3p1.0) specification. W3C proposed specification."
13. Cranor, L. F., M. Langheinrich and M Marchiori (2002). "A P3P Preference Exchange Language 1.0 (APPEL 1.0). W3C Working Draft."
14. Lederer, S., A. K. Dey and J. Mankoff (2002). "Everyday Privacy in Ubiquitous Computing Environments." Ubicomp 2002 Workshop on Socially-informed Design of Privacy-enhancing Solutions in Ubiquitous Computing.
15. Schreck, J. (2003). *Security and Privacy in User Modeling*. Kluwer Academic Publishers.
16. Westin, A. F. (1970). Privacy and Freedom.
17. Menezes, A. J., P. C. v. Oorschot and S. A. Vanstone (2001). *Handbook of Applied Cryptography*. CRC Press.
18. Python Library Reference (2004).
19. Nielsen, J. (1994). *Usability Engineering*. Morgan Kaufmann
20. Whitten, A. and J. D. Tygar. Why Jonny Can't Encrypt: A Usability Evaluation of PGP 5.0. 8<sup>th</sup> USENIX Security Symposium.
21. Fu, K., E. Sit, K. Smith and N. Feamster (2001). Dos and Don'ts of Client Authentication on the Web. 10<sup>th</sup> USENIX Security Symposium. Washington D.C., USA
22. OWASP (2002). A Guide to Building Secure Web Applications: The Open Web Application Security Project. 2004.
23. Brar, A. and J. Kay (2004). *Privacy and Security in Ubiquitous Personalized Applications*. Technical Report 561. School of Information Technologies, University of Sydney.

# A Single Sign-On Identity Management System Without a Trusted Third Party

Brian Richardson and Jim Greer

University of Saskatchewan, Department of Computer Science,  
ARIES Laboratory, Saskatoon, Saskatchewan, Canada  
{Brian.Richardson, Jim.Greer}@usask.ca

**Abstract.** Single sign-on identity management systems that rely on the use of a trusted third party, such as .NET Passport, face privacy and security risks when dealing with the management and storage of users' personal information. Presented in this paper is the design of a single sign-on system, which does not rely on the use of a third-party to manage users' personal information but at the same time allows users to present themselves to online businesses with more than one identity.

## 1 Introduction

Personalization of on-line content by on-line businesses can improve a user's experience and increase a business's chance of making a sale, but with stricter privacy legislation and Internet users' increasing concerns about privacy, businesses need to ensure they do not violate laws or frighten away potential customers. This paper describes the design of our Identity Management Architecture (IMA). The IMA system allows users to decide, on a per business basis what personal information is disclosed. It gives users greater control over their personal information held by on-line businesses, and does not rely on a trusted third-party for management of personal information.

In order to complete any commercial transaction on-line, people must provide personal information such as their name, address, phone number, email, credit card number, etc. On-line businesses often record more information than is actually needed to process a transaction. Some businesses monitor and record what types of products are bought or even the customer's browsing patterns. This is done to form a detailed profile that will allow the business to target a customer with future advertising of products more closely related to individual interests. As a result, businesses may be inadvertently in violation of privacy law and customers may be unaware of the extent to which their personal data is being stored or used. Individuals sometimes try to counter such actions by supplying false or misleading data in an attempt to conceal their identities. Thus, the following three factors formed the basis of this research:

1. Legislation: Canada's Personal Information Protection and Electronic Documents Act (PIPEDA) and how businesses can readily comply with this law [10].

2. Personal Concerns: The increasing concerns of Internet users about what information on-line businesses record about them.
3. Tool Support: The lack of an available privacy tool that allows for management of multiple identities.

Currently every popular personal information management system requires a third-party or a business to pass a user's personal information to another business. We believe that a personal information management system can be designed which does not rely on a third-party, and provides users with flexibility and control over the management of their personal information, while supporting business compliance with privacy legislation (such as PIPEDA). We also believe that privacy can be supported through the use of multiple identities by allowing a user to partition personal information into multiple pieces, each referred to as an identity. This would allow the user to choose on a per-business basis which identity to present. The use of multiple identities is a feature that is not yet offered by any commonly used personal information management system.

A Single Sign-On (SSO) system is one type of identity management system that allows a user to login to a system and gain access to numerous resources all with the use of a single username and password [14]. Systems like .NET Passport and Liberty Alliance allow for the use of a single username and password at multiple web sites, where the information for that one account can follow the user from site to site without the user having to re-enter information at each site visited, as long as that site is a participating member of that SSO service. These systems rely on some third party management of a user's personal information.

One common factor in most SSO systems is that they make the assumption that users always wish to present themselves online as the same identity with the same personal information. In fact people may not want all of their activities online to be linked to the same identity. Many people who use the internet will present themselves in with different identities when their purpose for using the internet changes. Each of these identities may contain some unique personal information with some overlap, and the owner of these identities would not want the activities taken while under each identity to be linked together. For example, if someone was considering finding a new job, he or she may not want the job hunt activities to be linked to a work identity for fear of his or her current employer finding out. It is for reasons like this that SSO systems need to realize the need for supporting multiple identities.

## **2 Background**

Two SSO identity management systems have emerged in the consumer e-commerce arena. The Passport system ([www.passport.net](http://www.passport.net)) was founded in 1999 by Microsoft. This system was designed to provide a single sign-in service that would allow Internet users to have one account for access to all Passport participating web sites [8]. The Passport system handles authentication of users by having the sign-in page on each participating web site authenticate the user by contacting the Passport system [8]. The Liberty Alliance Project ([www.projectliberty.org](http://www.projectliberty.org)) started in 2001 by Sun Microsystems to create an open standard, single sign-on authentication service [2].

The Liberty Alliance project has gained support from many well known organizations and now has more than 30 companies (e.g., Computer Associates, Hewlett Packard, Novell, etc.) involved in the development of the specification [2].

One of the main features of the Passport system is the single sign-in service. Although this is convenient, it does not allow for any sort of management of multiple identities. Passport does not allow the creation of multiple identities (i.e., more than one set of personal information) to be associated with a single account to allow a user to choose on a per business basis, what personal information a business receives. Although it is true that this could be accomplished by creating multiple Passport accounts, this defeats the purpose of a single sign-in service, since this approach would require a user to manage several Passport accounts, to remember multiple usernames and passwords, and to remember which account had been used at each business.

There is a misconception that Liberty Alliance is a similar service to .NET Passport. This is not the case. While .NET Passport is a single sign-on service implemented by Microsoft and used by online businesses, the Liberty Alliance Project is the development of a specification that can be implemented by businesses who wish to participate [11]. This specification allows businesses to form identity sharing relationships between each other and each implementation of the specification allows for this communication.

Liberty Alliance is based on the idea of allowing users to connect multiple sets of personal information, that exist across several on-line businesses, into one easy to manage identity. This allows for the convenience of a single sign-on service, as well as easier management of personal information across multiple businesses [1]. The Liberty Alliance architecture allows an Internet user to store his or her personal information with a trusted business. When the user needs to access a service provided by another business that is part of the same alliance of associated businesses, the user's chosen trusted business provides authentication of the user and forwards the user's identity information [1].

This system architecture is unique. Rather than relying on a trusted third party system, such as .NET Passport to provide a user's identity to each business the user accesses, it allows the user to have a business he or she trusts store and pass identity information from one business to another, which is part of the same group of associated businesses [1]. A group of associated businesses who have an agreement to share user identities and act as a single sign-on service is referred to as a Circle of Trust (COT). In a COT one business may act as the identity provider for a user and provide that identity to other businesses in the COT the user accesses [11]. One downside to this design is that identity management across multiple businesses is restricted to the set of businesses that have formed associations with each other. If a business is part of another group of associated businesses, identity information passing between these businesses is not available.

In early 2005, Microsoft announced that it was no longer going to pursue Passport as a solution for businesses to allow customers to manage their credit card and other personal information as they move from business to business online [6]. Companies such as eBay and Monster.com, two of the biggest non Microsoft companies to be using .NET Passport, who were initial supporters of the system, have dropped it in the last year. This, along with a lack of interest by many businesses and a failure of

internet users to openly accept Passport for storing their personal information, led to this decision by Microsoft.

There were concerns by companies about having Microsoft as the middle business between a company and its customers [6]. Users have been skeptical about storing their personal information in the Passport system. Anytime a user logs in to a Passport participating site, that site is immediately able to access all information in the user's Passport account [4]. Unfortunately for Microsoft, its systems often are the target of attacks by hackers, which has caused some embarrassment for Microsoft when security holes have been made public. A report by AT&T labs exposed several security flaws with Passport. One such flaw was that in order to compromise user accounts one only required a site that had a fake Passport login. This allowed usernames and passwords to be obtained providing access to all information in Passport about that user [4]. Microsoft also had to discontinue its use of the .NET Passport Wallet, which is a service that stored a user's credit card information, after it was discovered that all it would take to steal a person's Passport account and gain access to the Passport Wallet would be to get a user to open a Hotmail email [5]. Issues like this have raised continued concerns about security and privacy in the Passport system.

### **3 A Single Sign-On System Without the Third Party**

In order to understand what makes our Identity Management Architecture (IMA) system different from existing SSO systems, it is important to understand the overall architecture of the system.


The IMA system is designed around two main components: the IMA Manager, which is the client application, and the IMA Web Service, which is the web service deployed by participating on-line businesses. Each business that wishes to participate in the IMA network must follow the standard defined for the IMA Web Service, must implement this service, and deploy it on its web site. Through this service, all interactions with the IMA Manager Client application are handled. Each user who wishes to benefit from the IMA network installs the IMA Manager application on his or her computer. This application ties into the user's web browser. Each person using the IMA Manager can create one or more identities, and then when visiting the web sites of participating businesses, choose which identity to associate with each business. For future visits, this identity will be used by default unless a different identity is selected by the user.

One of the key features of the IMA system, which is not offered by other personal information management tools, is the ability to create and manage multiple identities from within a single user account. The use of multiple identities does more than just restrict what information a business will see, but also allows people to interact with a business for more than one purpose. For example, if someone shops at an on-line computer parts store, sometimes for work purposes and other times for personal purposes, he or she may want to create a "Work" identity and a "Personal" identity. Creating separate identities allows someone to more easily manage these two separate relationships with a business. This may be beneficial to a business, especially one that personalizes web site content based on the interests of the user. If someone is using

his or her “Personal” identity, the business may use the browsed products and recently purchased products to make suggestions to the user about other products that may be of interest. If this same person visits the business’s web site at another time using the “Work” identity, the business will be better able to tailor content towards the interests associated with this identity.

One goal of the Identity Management Architecture is to avoid use of a trusted third party system and to not require businesses to communicate with each other for the purpose of providing a customer’s information. The IMA system has two main components:

1. IMA Manager (Client): An application that attaches to the user’s web browser and handles the management of all user identities and web browsing history.
2. IMA Web Service (Business): A web service that each participating business provides to allow users of the IMA Manager to send and receive identity information.



**Fig. 1.** The overall architecture of the IMA system: an IMA client application connects with a participating business using the IMA web service.

The IMA Manager allows a user to contact a business’s IMA Web Service to make a request to see what information the business currently has stored in the user’s profile. A user may correct or remove information. If the information to be changed is contained in an identity, a user may modify the identity information stored in the IMA Manager and the application will automatically forward updates to all businesses associated with this identity. The user may associate another identity with a business at any time; this will be used for future visits to that business.

As shown in Figure 1, the IMA Manager runs on the user’s local web browser. The Manager receives from the web browser the URL of each site the user visits while on-line. It then checks to see if the business is participating in IMA by attempting to contact the IMA web service that all participating sites are required to make available. If this business is not participating, then this is shown in the IMA Manager’s display. However, if the web service is available, this URL is stored and the service is contacted. The first communication with a business is the transmission of the user’s preferred on-line identity. From this point on, each time the user returns to this web site the user will be identified by this identity, allowing the business to asso-

ciate information, such as the products browsed, with this identity in order to determine the user's interests.

Only that information contained in a single identity is sent to a business. This is unlike Passport, which uses only one set of personal information for a user's account and provides it to each partner business. If a user wishes to have more than one identity, multiple Passport accounts must be created which defeats the purpose of a single sign-in service. The IMA system is designed to provide management of multiple identities on behalf of the user.

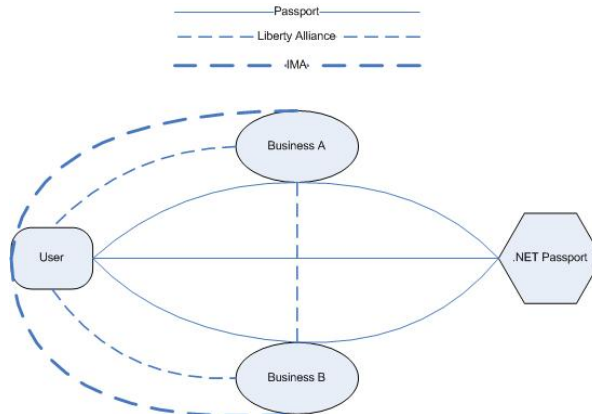
One additional key feature of the IMA system is that it provides users with the option to directly request what information a business has about them. The only information users have access to view and update with Passport is the personal information that was entered when the account was created. This does not provide users with any access to additional information a business has about them and it does not provide any benefit to businesses in terms of compliance with information disclosure requirements placed on businesses by privacy legislation.

With most personal information management systems (single sign-on services) it is necessary for either the user or service provider (or both) to establish additional relationships (i.e., with a third party system or another business) in order to be able to participate. The main goal of the IMA system is to provide the same types of services provided by traditional personal information management systems, but without requiring additional relationships to be established.

In order for users to participate in the .NET Passport system they must create an account with .NET Passport and provide all personal information they wish to be used in that account. This is the first new relationship that must be established outside of the traditional customer-business relationship. The second new relationship required is between .NET Passport and each business that wishes to participate (see Figure 2). These two new relationships require both users and businesses to be willing to participate in a personal information management system that involves a trusted third party.

The Liberty Alliance system allows users to select a business they trust to store their personal information. In order for this identity to be used at another business two conditions must be true: the business users are visiting must also be participating in the Liberty Alliance system and both businesses must have established an identity sharing relationship with each other (see Figure 2). If both are true, then users may use their accounts at other businesses and have their personal information transferred from the trusted business.

The IMA system relies on the existing customer-business relationship. If both the user and business are participating in the IMA system, personal information the user has in an account can be transferred to each business by the user (see Figure 2). The IMA system does not require a user or business to have to form additional relationships with either a third party system or another business.



**Fig. 2.** Shown in this figure are the relationships that exist in .NET Passport, Liberty Alliance and the IMA network. As shown in the figure, Passport requires a user and business connection to a third party service, Liberty Alliance requires a business to business connection, while the IMA network requires neither relationship.

## 4 IMA System Implementation

In order to demonstrate the design of this system, a prototype implementation of IMA has been built. The implementation includes the client application as well as a sample implementation of the services associated with a typical participating business. Each component of the system has been written using the .NET framework [7]. Four components were implemented: IMA Toolbar, IMA Manager, IMA Web Service, and an example participating online business.

The IMA Toolbar is a .NET application that integrates into the user's web browser and provides the user with information on the participation of a web site being visited and the identity currently being used. This toolbar allows the IMA Manager to be automatically started when the user opens a web browser. This application is a class library that is registered as a toolbar with Internet Explorer and controls the IMA Manager. The IMA Toolbar created for this project was based on an example .NET Toolbar obtained from The Code Project [15].

The IMA Manager is a .NET application that runs in the background on the user's system. It is displayed as a taskbar icon, unless the user wishes to view the browsing history, or modify identities. This application is a standard windows application that displays a window when the taskbar icon is clicked.

The IMA Web Service is a .NET web service that allows the IMA Manager application to communicate with a participating business. Identity information is transferred to and from the IMA Web Service as an XML document. All other information recorded by the business that has been associated with the current identity can be retrieved in an XML document that the IMA Manager can display to the user who can make changes if necessary. How a business actually implements this system and how it ties into its database is the decision of the business. All that is required is that the

URL of the service and the methods offered match the ones required by the IMA Manager

To demonstrate the IMA system a small demo E-Commerce web site was built using ASP .NET. This site acted as a participating business providing the IMA Web Service that allowed the IMA Manager to be used to experiment with how the IMA system would work. This allowed for a better understanding of the IMA system design and also played a role in several design changes. The design of the IMA Manager and the IMA Web Service are described in detail in the following sections.

The IMA system allows users to set different identities, for example for personal, work, and private related activities. Having the ability to easily separate these identities allows web users to ensure that the information recorded about them at certain web sites is based on the identity they have assigned for that site. Users are able to select a single identity to be used for all web sites they visit, but they are able to switch to another identity when desired. The IMA Manager also allows users to view a list of recently visited business sites and modify defaults in order to have some alternate identity associated with any site at any time. This change results in the information for the alternate identity being immediately forwarded to the business. This alternate identity is then used for subsequent visits to this site.

When shopping at a number of different on-line businesses, it is convenient for customers to have each business keep some of their personal “account” information such as name and mailing address etc., and to connect to that account with a username/password. However, users are normally required to update account information manually at each business, when personal “account” information changes. The IMA Manager allows a user to update an identity and have those changes automatically forwarded to all businesses with which this identity has been associated.

If a business is participating in the IMA system, there may be no need for a first time visitor to the web site to have to fill out a long form to access a feature the business offers, as is often required for download of trial software. Instead, as the user accesses the site, he or she may choose to associate an identity with the business. Only after the IMA Manager receives confirmation from the user is the identity forwarded to the business by the IMA Manager.

#### **4.1 Limitations**

Personal information in the IMA system is always stored in two places; at the participating businesses the user has visited, and on the user’s computer. Any responsible business will have security measures in place to protect stored personal information. However, the information stored on an individual’s computer may be at risk of being compromised. In order to ensure no one other than the owner of the account has access to this information the account is stored in an encrypted, password-protected file. The user sets the username and password for an account when it is created. No one other than the owner of the account knows the password. Since there is no third party to store the account information or password, if the user forgets the password there is no way to retrieve it.

Even if an IMA participating business has excellent security measures in place to ensure each user's personal information is protected, there may still be increased security risks since this architecture promotes a more open exchange of personal information between users and businesses. If someone tried to make a request to a business for personal information while posing as another user, this could lead to the business disclosing a user's personal information to the wrong person. The way the IMA system attempts to reduce this risk is through the use of unique keys stored in each identity. When an identity is created, a unique key is generated and added to the identity. This unique key is used by businesses to authenticate an identity each time it is used. The key is never known explicitly by the user, but instead remains in the identity of the user account and is provided to businesses along with the other client-side information in the identity. In order for someone to pose as another user to retrieve identity information from a business, the impostor would have to know the key for that user's identity or be using the client computer of that user.

Another potential threat to the IMA system is that a disreputable business may not publicly state that it is participating yet may secretly deploy an IMA Web Service. Through this service the business may attempt to extract personal information from visitors to the business's web site if those users are using the IMA Manager. The result is that a business may try to collect personal information from users, yet not provide the required access for users to stored personal information. For the IMA Manager to release an identity to a business this action must be explicitly initiated by the user. First the user must be currently visiting the business's web site. Second the user must select an identity and attempt to associate it with the business. Third the user is asked by the IMA Manager to confirm the sending of the identity. The IMA Manager will never automatically release information to a business. If the user confirms that they want the identity sent to the business, then only at this time is this action taken by the IMA Manager. The IMA Manager makes every effort to prevent a business from receiving identity information without full knowledge of the user.

## 5 Related Work

There is currently extensive research work going on in the area of identity management by the Privacy and Identity Management for Europe (PRIME) [9] and The Future of Identity in the Information Society (FIDIS) [3] projects. Work by both of these projects has produced prototype identity management systems which each take a similar approach to handling of identities that allows users to switch identities (roles) based on which identity they wish to present to an organization. While both the PRIME [12] and FIDIS [13] projects are large in scope and are exploring in great depth many privacy and identity management issues, the IMA system really only attempts to address a small portion of the identity management problem.

The FIDIS project presents the prototype iManager which is the Identity Manager for Partial Identities [13]. Each partial identity contains a subset of the user's information that is applicable to the information needed for the user's current role, such as an identity that contains a credit card number and mailing address used when the user is shopping online. This approach is similar to how identities are handled in the IMA

system. Basically each identity in IMA, like partial identities, is identified and authenticated by a unique key that allows the user to be authenticated by an organization for all future visits using the same identity. This allows an organization to be able to track all repeat visits and associate information with that identity about the user, allowing the user to build up a relationship, regardless of whether or not the user has even provided his or her identifying personal information such as name, address, email, etc. A similar approach to the iManager's partial identities was followed in the design of the IMA system's multiple identity user account. The IMA system splits the user account up into identities which each contain a different subset of the user's personal information, identities such as anonymous, personal, work, school, etc. Each identity allows the user to only provide that set of information contained within the given identity, all other identities and information contained in the account are not disclosed to an organization.

The PRIME project presents the prototype IDM system [12]. The IDM system is a much more overall solution than the IMA system, however there are still some areas of IDM that the IMA system touches on. While the IDM system improves privacy by allowing the user to remain anonymous, even during a transaction, assuming there is a trusted third party that in the case of a problem (e.g., legal matter) the identity can be recovered, however in the IMA system no attempt to preserve anonymity like this is made. In the IMA system if a user decides to complete a transaction with a business, it is up to the user to decide whether or not he or she wishes to disclose an identity containing the required information, no anonymity is preserved in this type of transaction.

The primary goal of the IMA system was to build a single sign-on system that did not require third party storage or knowledge of user's information. As an additional feature the IMA system would also allow users to create and manage more than one identity from within a single user account where all identities could be accessed by a single username and password. These initial requirements were what the design of the IMA system had to achieve. Rather than comparing the design of the IMA system to the ongoing work in identity management taking place in PRIME and FIDIS, the IMA system was looked at more as an improvement upon existing single sign-on systems such as .NET Passport and Liberty Alliance. Both of these existing systems have well defined architectures which allow for a more detailed comparison to be made.

## **6 Future Work and Conclusions**

Our next steps will be to build a larger, more complete implementation of the IMA system. The first issue that will need to be addressed is security. This will include determining the most secure way for a business to identify an IMA system user, and ensuring that only businesses authorized to receive the user's information have access to it. The second issue that will need to be decided is how to allow users access to their accounts from multiple locations. Since the IMA system does not rely on a third party system for account access and storage, the IMA system will require a different approach. There are several options available. One would be to store an encrypted

account file with a central service (whose trust level would be far reduced from that of the Passport approach). Such an approach allows for anonymous and secure storage of personal information in a "safety deposit box" that may be retrieved and used from various locations. A final decision on this matter has not been made.

The IMA system provides a design for a personal information management architecture that is offered as an alternative to .NET Passport and Liberty Alliance.

The main purposes of this research work are to demonstrate the benefits to a single sign-on system of increased access for users to their personal information and the allowing of users to maintain more than one identity. The main contribution of this project has been the design and development of an architecture for an identity management system that does not use a third-party or require businesses to transfer identity information from one business to another. It is hoped that this work will be the basis for more research into identity management systems (i.e., single sign-on systems) that provide users with more control over who can view or use their personal information, while allowing businesses to increase their compliance with privacy legislation, and improve the privacy of Internet users.

## References

1. S. Cantor et al., "Liberty ID-FF Architecture Overview" 2003;  
<http://www.projectliberty.org/specs/liberty-idff-arch-overview-v1.2.pdf>.
2. J. Evers, "EBay Cancels Its Passport" Jan. 2005;  
<http://www.pcworld.com/news/article/0,aid,119137,00.asp>.
3. FIDIS, "Future of Identity in the Information Society"; <http://www.fidis.net/>.
4. S. Johnston, "Pondering Passport: Do You Trust Microsoft With Your Data?" Sept. 2001;  
<http://www.pcworld.com/news/article/0,aid,63244,00.asp>.
5. B. McWilliams, "Stealing MS Passport's Wallet" Nov. 2001;  
<http://www.wired.com/news/print/0,1294,48105,00.html>.
6. J. Menn, "Microsoft's Passport fails to travel far as Web strategy" Dec. 2004;  
[http://seattletimes.nwsource.com/html/business/technology/2002136272\\_passport31.html](http://seattletimes.nwsource.com/html/business/technology/2002136272_passport31.html).
7. Microsoft, "Microsoft .NET Framework Developer Center";  
<http://msdn.microsoft.com/netframework/>.
8. Microsoft, "Microsoft .NET Passport Privacy Statement" 2003;  
<http://www.projectliberty.org/specs/liberty-idff-arch-overview-v1.2.pdf>.
9. PRIME Project, "Privacy and Identity Management for Europe"; <http://www.prime-project.eu.org/>.
10. Privacy Commissioner of Canada, "A Guide for Individuals" Nov. 2003;  
[http://www.privcom.gc.ca/information/02\\_05\\_d\\_08\\_e.asp](http://www.privcom.gc.ca/information/02_05_d_08_e.asp).
11. P. Roberts, "Liberty Alliance Explains Its Sign-On Services" Feb. 2003;  
<http://www.pcworld.com/news/article/0,aid,109277,00.asp>.
12. WP 14.1, "Framework V1" Mar. 2005; [http://www.prime-project.eu.org/public/prime\\_products/deliverables/fmwk/pub\\_del\\_D14.1.a\\_ec\\_wp14.1\\_v1\\_final.pdf](http://www.prime-project.eu.org/public/prime_products/deliverables/fmwk/pub_del_D14.1.a_ec_wp14.1_v1_final.pdf).
13. WP3, "Structured Overview on Prototypes and Concepts of Identity Management Systems" 2004; [http://www.fidis.net/fileadmin/fidis/deliverables/fidis-wp3-del3.1.overview\\_on\\_IMS.pdf](http://www.fidis.net/fileadmin/fidis/deliverables/fidis-wp3-del3.1.overview_on_IMS.pdf).

- 14.R. Yasin, "What is Identity Management?" Apr. 2002;  
[http://infosecuritymag.techtarget.com/2002/apr/cover\\_casestudy.shtml](http://infosecuritymag.techtarget.com/2002/apr/cover_casestudy.shtml).
- 15.P. Zolnikov, "Extending Explorer with Band Objects using .NET and Windows Forms" Apr. 2002; <http://www.codeproject.com/csharp/dotnetbandobjects.asp?print=true>.

# Intra-Application Partitioning of Personal Data

Katrin Borcea, Hilko Donker, Elke Franz, Katja Liesebach,  
Andreas Pfitzmann, and Hagen Wahrig

Dresden University of Technology, Dresden, Germany  
{borcea|donker|ef1|liesebach|pfitza|wahrig}@inf.tu-dresden.de

**Abstract.** Personalization provides users a comfortable working environment. But the necessary collection of personal data can imply privacy problems. Usual approaches to minimize privacy problems aim at separating data disclosed in different applications. However, this inter-application partitioning is not sufficient in case of large applications. Here we introduce the concept of *intra-application partitioning* of personal data by means of application-internal contexts. The description of such task-related contexts enables users to assess the linkability of their actions within the application. This approach helps users to control by themselves the linkability of their personal data.

## 1 Introduction

Currently, most applications aim at providing users a comfortable working environment adapted to personal needs. This goal requires to collect and to evaluate personal data describing, e.g. users' preferences and goals. Any collection of personal data, however, puts privacy at risk. Since computing systems are not perfectly secure, we cannot exclude unintended access to the data.

Hence, as less data as possible should be collected and, if possible, no personal data at all should be disclosed. Privacy-enhancing Identity Management (PIM) realizes decentralized data management and transparent data processing for users [1]. The users are enabled to partition their personal data and to control disclosure of data subsets [2, 3]. Each subset of information is called a partial identity [5], which must be unlinkable by others except their owner. Therefore, uncorrelated pseudonyms are used as identifiers for the partial identities.

Usually, PIM is used to keep data in different applications separate from each other (inter-application partitioning), or to separate actions of a user within different roles. Particularly, for applications comprising different services and/or interactions with other users, the possibility to partition personal data is needed - even for repeating actions. Within this paper, we introduce possible solutions for this *intra-application partitioning* depending on *application contexts*.

In contrast to inter-application partitioning, for intra-application partitioning the application must be aware of the partitioning and must support it. Therefore, the application needs to be modified. Additionally, we have to consider interactions between users, i.e. former actions of the user under a pseudonym, and interactions with other users (possibly) under different pseudonyms.

For the realization of intra-application partitioning, we aim at utilizing mechanisms provided by the PIM architecture currently being developed within the EU project PRIME<sup>1</sup>. As an example application that will be executed on top of this PIM architecture [1], we will use the eLearning application BluES<sup>2</sup>. It is based on a client-server architecture and comprises conflicting requirements: Personalization is important for providing a reasonable working environment. Otherwise, severe privacy threats may imply acceptance problems of eLearning.

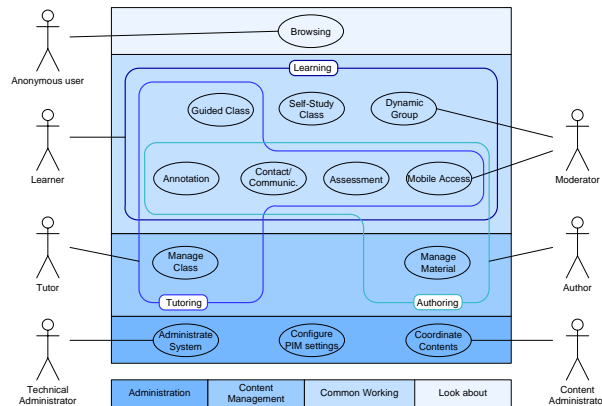
There are some approaches known in the field of eLearning which aim at enabling users to control disclosure of personal data. The authors of [9] use privacy policies to describe which data may be disclosed to whom and under which conditions. However, the partitioning of personal data is not considered. The approach introduced in [4] considers the use of PIM for using different pseudonyms for different services, but it neglects partitioning within one application.

In Sec. 2, we give a short overview of our example application. Afterwards, we introduce our definitions for contexts and general aspects in Sec. 3. In Sec. 4 we discuss possible realizations. Finally, Sec. 5 concludes and gives an outlook.

## 2 Short Overview on our Example Application

### 2.1 General Overview on BluES Concepts

eLearning comprises a number of use cases. Users can act within different roles in the eLearning environment. Within this paper, we refer to the use cases provided by BluES [1] which support *learning*, *authoring*, and *tutoring* (Fig. 1).



**Fig. 1.** Overview on use cases within the eLearning application BluES.

<sup>1</sup> <http://www.prime-project.eu.org>

<sup>2</sup> Java based eLearning environment, <http://www.blues-portal.de>

Our use cases can be classified into four categories: administration, content management, common working, and look about (anonymous browsing). BluES supports different learning scenarios: *guided class* (i.e. assistance by tutors), *self-study class*, and *dynamic groups* (i.e. support of cooperative learning). Common working provides additionally functionality, e.g. annotations and communication. Nearly all of the use cases shown in Fig. 1 comprise a number of sub use cases. Thus, a guided class contains, e.g. the sub use cases registration, process learning modules, practice, and examination<sup>3</sup>.

The additive use case *Configure PIM settings* enables users, e.g. to configure the partitioning of personal data and, thereby, to control the disclosure of personal data by themselves as required in Sec. 1.

## 2.2 Personalization and Privacy Requirements

The eLearning environment should automatically adapt the content presentation and generate hints depending on the learner's progress. This assistance implies acquiring of personal data such as learning history and his preferences.

Cooperative scenarios pose special requirements on personalization. If a dynamic group is established, users who are possibly also interested in the topics addressed in this group should get a hint from the system. Therefore, the application needs information about the interests of the users for this support.

The application as well as assisting tutors can provide the best support if they to recognize the users and if they know as much as possible about them. Even if suitable user models reduce the amount of necessary information [8], collecting and evaluating of personal data is inevitable implying that actions of a user are linkable. This linkability facilitates to create detailed user profiles [1]. Finally, those profiles allow to draw conclusions about the user (e.g. about his habits, equipment, or social status) possibly leading to a biased environment.

Therefore, a fine-grained partitioning of personal data within the application itself is essential in order to enable unrestricted behavior of users.

## 3 Contexts and Context Switches

### 3.1 Basic Considerations

Within BluES, contexts are used as means to realize intra-application partitioning of personal data. The term **context** was first introduced in [7]: Computation does not occur at a single location in a single context but rather spans a multitude of situations and locations. We use this term in a similar way to describe a specific situation in which a user works to perform a task. The application must be able both to detect the current context and to determine which actions (based on this context information) the users will take. Context information is derived

---

<sup>3</sup> A more detailed description of the use cases and sub use cases can be found at <http://blues.inf.tu-dresden.de/concepts/useCases.html>.

from diverse information sources, such as user actions. A **context switch** means that the user starts working on another task.

We assume that differentiating contexts is intuitive to each user, since this simply means to differentiate tasks. If a user starts to work on a new task, he can scrutinize which other users will recognize him, what they possibly already know about him, and which information about himself he is willing to disclose to them. This implies the need for a privacy-aware user interface.

Possible context switches are predetermined by the functional structure of the application. Any action that does not belong to the current (sub) use case implies a context switch that is a potential point to change the pseudonym (switch of partial identity). Actual pseudonym switches depend on users' privacy desires as well as on application restrictions specified by means of policies.

Even if users decide to switch to another partial identity, we have to consider linkable partial identities. An observer who can see all pseudonyms at server side can at least conclude there might be pseudonyms which belong to one and the same user. Users acting at the same time in the same way establish an anonymity set [5]. The size of such anonymity sets restricts the achievable privacy: the larger the anonymity set – the smaller the degree (the probability) being identified.

### 3.2 A First Approach of Modeling Contexts

For an automatic evaluation of potential context switches, first, we need a description of context and context switch. A context  $C_i$  is described by a set of attributes. Two contexts must differ in at least one of these attributes. Generally, there are four classes of attributes: a general use case description  $uc_g$  (top level use case, e.g. guided class = "Advanced Statistics"), an additional use case description  $uc_a$  (sub use case(s), e.g. synchronous learning = "Hypothesis"), a history  $h$  describing actions performed within the corresponding sub use case (e.g. pose\_a\_question), and a description of communication partners  $cp$ :

$$\begin{aligned} C_i &= (uc_g, \{uc_a, \{h, \{cp\}^*\}^*\}) \\ uc_g &= (role, top\_level\_use\_case) \\ uc_a &= (sub\_use\_case\_chain) \\ h &= (date, time, action, \dots) \end{aligned}$$

A change of one of the attributes of  $uc_g$  implies a context switch in any case. If the user wants to pool different sub use cases within a top level use case,  $uc_a$  contains a list of the included sub use cases. If a user wants to separate sub use cases,  $uc_a$  contains only one sub use case. The finest possible granularity – *atomic contexts* – can be achieved if each action within a sub use case is separated. In that case,  $uc_g$ ,  $uc_a$ , and  $h$  include single values only.

This context description depends on the chosen granularity which attributes are evaluated in order to recognize a context switch. These attributes are called *differentiating attributes*, which comprise at least  $uc_g$ . They also include  $uc_a$ , if the user wants to separate different use cases, and  $h$ , if the user even wants to separate actions performed within a sub use case. Since we assume that the user

does not perform different actions at the same time, *cp* is not used to distinguish contexts. However, *cp* provides linkability information to the user.

### 3.3 Describing Context Switches

In fact, only actions of a user that initiate communication with the server require a decision whether a context switch is implied. We call these actions *visible actions* since they are visible outside the client machine.

In case of a visible action, the application evaluates the current context of the user and the invoked action. Since any invocation includes information about the corresponding use case, the application can check the differentiating attributes. We distinguish between the following cases:

- *Mandatory context switch*: The current context and the visible action differ in at least one of the attributes of  $uc_g$ .
- *Potential context switch*: The current context and the visible actions differ in at least one of the differentiating attributes, but not in an attribute of  $uc_g$ .
- *Actual context switch*: According to the result of the evaluation, the visible action actually implies a context switch. Mandatory context switches always result in an actual context switch.

We distinguish between potential and actual context switches since we have to consider different possibilities to realize context switches. The scope of a pseudonym can comprise one or more contexts. However, the suggested structure for the description of contexts provides a reasonable basis for partitioning.

Due to the monotony of anonymity (confidentiality goals can only decrease over time [6]), it is not possible to split scopes of pseudonyms *ex post*. The linking between the newly introduced partial identities is already known to the others. Thus, context switches within a top level use case are only necessary if the user wants to change his pseudonym at this point.

### 3.4 Handling Context Switches

Basically, the application has to use a general event framework. Each invocation of a function, i.e. an action, is handled as event. Visible actions initiate an evaluation of potential context switches. Potential context switches can

- directly imply an actual context switch,
- imply an actual context switch as well as a notification for the user about this context switch, or
- initiate an interaction with the user where the user has to decide by himself whether an actual context switch should be initiated.

Configurations as well as already created contexts are managed and stored solely at client side. The server is not able to link different partial identities which are generated within the application. This approach enables users to control which of their data can be linked together.

## 4 Configuring Potential Context Switches

### 4.1 Possible Approaches

Basically, we can distinguish three fundamental strategies to reasonably configure potential context switches: prospective, collateral, and retrospective. Each of these concepts is based on the definition of atomic contexts.

*Prospective configuration.* Before working with the application, the user defines own contexts and points to switch to another partial identity. Afterwards, the defined contexts are valid for the whole session. On basis of these settings the system switches the contexts as well as the partial identities without user notification while he is working with the application.

*Collateral configuration.* In contrast to the first approach, every visible action requires a one-time decision: The user has to decide if he wants to switch to another context as well as to another partial identity.

*Retrospective configuration.* Every visible action directly implies a context switch and a switch of the partial identity. The user decides retrospectively which of these contexts and partial identities shall be linked together. This approach delivers indications for reasonable context classifications as well as for supporting especially cooperative learning.

Table 1 discusses advantages and disadvantages of these approaches.

Configuration Approach	Advantages	Disadvantages
Perspective configuration	Users are not disturbed in their working and thinking processes by system queries concerning possible context switches.	Not suitable for unexpected and new situations; users should be very familiar with the application as well as with privacy aspects.
Collateral configuration	Offers the best flexibility; is particularly suitable for personal adaptations to current requirements.	Users are disturbed each time there is a possible context switch.
Retrospective configuration	Per default, this approach provides the best anonymity at run time.	The user must actively analyze his data track backwards in order to summarize possible contexts.

**Table 1.** Advantages and disadvantages of the different configuration approaches.

To conclude, none of the strategies can be used stand-alone. Only a combination of the three concepts is reasonable. Particularly, an eLearning environment is too complex for perspective configuration. Unknown situations must always imply interactions with the user. Furthermore, there should be the possibility to analyze and to link contexts as well as partial identities backwards in order to derive strategies for future decisions.

Fig. 2 gives an overview of the three different concepts. The example scenario illustrates that the user has access to his personal workspace in the scope

of pseudonym "Ps 1". Further, he attends the self-study class "Principles of Statistics" and the guided class "Advanced Statistics". In the latter, he separates different actions using the pseudonyms "Ps 2" and "Ps 3".

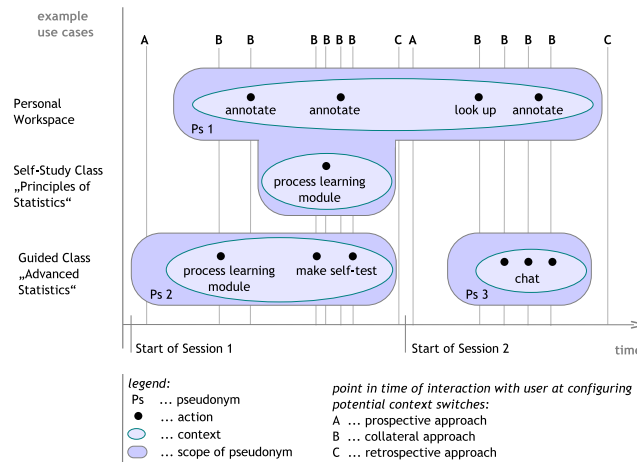


Fig. 2. Three approaches of configuring potential context switches.

## 4.2 Supporting Users in Defining Context Switches

Even if we assume that differentiating tasks is intuitive for users, configuring potential context switches might be not easy for them, especially, if users are not yet familiar with the application or with privacy issues. Therefore, support for users is an essential criteria for the acceptance of intra-application partitioning.

Thus, reasonable pre-configurations for context switches should be provided (context switch with/without notification, or interaction). Users could decide to download such privacy settings from a trusted provider and integrate into their application. When the users become more familiar with the application structure and with data partitioning they are enabled to adapt the default settings to their actual needs. The configurations should be possible for types of use cases as well as for specific use case instances. For example, users can define rules for guided classes but overwrite these general rules for a specific guided class they attend.

Of course, the user interface of the application must be enhanced in order to display the user's current privacy status.

## 5 Summary and Outlook

We described possibilities for intra-application partitioning of user data. This approach enables users to preserve their privacy even in applications which support

interaction with other users. The application itself must enable this partitioning. Therefore, we characterized application contexts used for the partitioning. The application should support flexible configuration of potential context switches. In order to reasonably support users we suggested to deliver pre-configurations.

There are still many open issues: Assessing linkability requires thoroughly further investigations. Within this paper, we have only discussed contexts and context switches from the point of view of a single user. However, multilateral interactions in cooperative scenarios will pose further requirements that concern a set of users. Furthermore, we have to consider linkability due to observation of context switches.

Another topic of future work is visualization of linkability, i.e. developing a suitable user interface, that helps users to decide whether they should generate a new partial identity or select an existing one. This includes information about other users as well as information about the privacy state of the user himself.

Currently, we are realizing the approach described in this paper. Our goal is to perform test trials with a reasonable user set in order to investigate performance aspects as well as user acceptance and necessary effort.

We like to thank Mike Bergmann, Sebastian Clauß, and Thomas Kriegelstein for helpful discussions. The work reported in this paper was supported by the IST PRIME<sup>4</sup> project; however, it represents not necessarily the view of the project.

## References

1. K. Borcea, H. Donker, E. Franz, A. Pfitzmann, and H. Wahrig. Towards Privacy-Aware eLearning. Accepted for PET 2005.
2. S. Clauß and M. Köhntopp. Identity Management and its Support of Multilateral Security. *Computer Networks*, (37):205-219, 2001.
3. S. Clauß, A. Pfitzmann, M. Hansen, and E. V. Herreweghen. Privacy-Enhancing Identity Management. The IPTS Report 67, pages 8-16, Sept. 2002. <http://www.jrc.es/pages/iptsreport/vol67/english/IPT2E676.html>.
4. T. Klobučar, V. Seničar, and B. J. Blažič. Privacy and personalization in a smart space for learning. *Int. J. Cont. Engineering Education and Lifelong Learning*, 14(4/5):388-401, 2004.
5. A. Pfitzmann, M. Hansen: Anonymity, Unobservability, Pseudonymity, and Identity Management — A Proposal for Terminology. Draft status: [http://dud.inf.tu-dresden.de/literatur/Anon-Terminology\\_v0.21.pdf](http://dud.inf.tu-dresden.de/literatur/Anon-Terminology_v0.21.pdf); September 2004.
6. A. Pfitzmann and G. Wolf. Properties of Protection Goals and their Integration into a User Interface. *Computer Networks*, 32:685-699, 2000.
7. B. N. Schilit, N. Adams, and R. Want. Context-Aware Computing Applications. In *Proc. of the 1st Int. Workshop on Mobile Computing Systems and Applications*, pages 85-90. IEEE, 1999.
8. J. Self. Bypassing the intractable problem of student modelling, 1988.
9. G. Yee and L. Korba. Privacy Policies and their Negotiation in Distance Education. In P. Derbyshire, editor, *Instructional Technologies: Cognitive Aspects of Online Programs*. IRM press, 2004.

---

<sup>4</sup> The PRIME project receives research funding from the European Community's 6th Framework Programme and the Swiss Federal Office for Education and Science.

# Privacy-Enhanced Collaborative Filtering

Shlomo Berkovsky<sup>1</sup>, Yaniv Eytani<sup>1</sup>, Tsvi Kuflik<sup>2</sup>, Francesco Ricci<sup>3</sup>

<sup>1</sup> Computer Science Department, University of Haifa, Israel  
{slavax, ieytani}@cs.haifa.ac.il

<sup>2</sup> Management Information Systems Department, University of Haifa, Israel  
tsvikak@is.haifa.ac.il

<sup>3</sup> ITC-irst, Trento, Italy  
ricci@itc.itc

**Abstract.** Current implementations of the Collaborative Filtering (CF) algorithm are mostly centralized and the information about users (their profiles) is stored in a single server. Centralized storage poses a severe privacy hazard, since user profiles are fully under the control of the recommendation service providers. These profiles are available to other users upon request and are transferred over the network. Recent works proposed to improve the scalability of CF by distributing the stored profiles between several repositories. In this work we investigate how a decentralized approach to users' profiles storage could mitigate some of the privacy concerns of CF. The privacy hazards are resolved by storing the users' profiles only on the client-side so they are used for computation similarity only on the client-side. Only a value indicating the similarity is transferred over the network, without revealing the profile itself. To further avoid the disclosure of the user's profile through a series of attacks, we propose that the users hide or obfuscate parts of their profile. Experimental results show that relatively large parts of the user's profile could be obfuscated without hampering the accuracy of the CF.

## 1 Introduction

Collaborative Filtering (CF) is commonly used in the E-Commerce realm for producing recommendation for various products. CF is based on the assumption that people with similar tastes prefer the same items. In order to generate a recommendation, CF initially creates a neighborhood of users with the highest similarity to the user whose preferences are to be predicted. Then, it generates a prediction by calculating a normalized and weighted average of the ratings of the users in the neighborhood.

In CF, user profile is a feature-vector containing information about user preferences with respect to a set of item the user rated. For quite some time CF has been applied in E-Commerce and direct recommendations of various kinds [15]. Personalized information delivery in general and purchase recommendations (that applies collaborative filtering) in particular can increase the likelihood of a customer making a purchase, compared to non-personalized approaches.

However, personalization brings with it the issue of privacy. Privacy is an important challenge facing the growth of Internet and the acceptance of various transaction models supported by Internet. Basically, Web users leave identifiable tracks while surfing the Web and there is a growing awareness and concern about the misuse of such information [1]. Many eavesdroppers on the Web violate users' privacy for their own commercial benefits, and as a result, users concerned about their privacy refrain from using useful Web services to prevent exposure [6], [4].

The need for protection of Web users' privacy triggers growing research efforts nowadays. Wide variety of research directions for preservation the privacy while surfing the Web in general, and for CF in particular are explored. Canny [2], [3] suggests privacy preservation approach based on peer-to peer techniques. He suggests forming users' communities, where the community will have an aggregate user profile, representing the group as whole and not individual users. Personal information will be encrypted and the communication will be between individual users and not servers. Thus, the recommendation will be generated on the client side.

Polat and Du [11], [12] suggest another method for preservation of user privacy on the central server by adding uncertainty to the data by using a randomized perturbation technique. Hence, the server (or the data collector) has no knowledge about true values of individual users rated items. They demonstrate that this method does not lower considerably the obtained accuracy of the results

In today's dynamic environments, formation of a community of users poses limitation on possibilities of information sharing. Users looking for personal information in various domains and situations may need to interact with different set of users every time. To accomplish this while preserving the users' privacy, we propose an approach that combines benefits from both Canny, and Polat and Du. Generally, we believe that users should control when and what personal information they would like to reveal.

Individual users may participate in a virtual, distributed CF system in the following way: every user may keep and maintain his personal profile of rated items. Recommendation is requested by a user sending parts of his profile and a request for recommendation. Other users may decide to respond to that request by sending their recommendation and degree of similarity with the requester. The originator collects the responses and uses them as a source for neighborhood formation that later on leads to local generation of recommendation. Such approach preserves users' privacy by leaving them in control of their personal information, while allowing them to support recommendation generation by other users.

We experimented with the publicly available Jester dataset containing rating for 100 jokes [5]. Results clearly demonstrate that adding the proposed privacy enhancements do not severely affect the accuracy of the recommendations obtained basing on the CF algorithm.

The rest of the paper is structured as follows: Section 2 discusses limitations of centralized CF and distributed CF. Section 3 presents our "on-the-fly" CF approach for recommendation generation. Section 4 presents the experimental results validating the approach and discusses open research questions. Section 5 concludes the work and presents directions of future research.

## 2 Distributed Collaborative Filtering and Privacy

Centralized CF systems have a number of disadvantages. In particular, they pose a severe threat to users' privacy, as the service providers collect valuable information about the users. This information can be transferred or sold once it was collected and used of malicious purposes. Thus, according to a recent survey [4], most users will not agree to divulge their private information. This causes users to refrain from providing personal information or to provide false information. Using CF without compromising user's privacy is most certainly an important and challenging issue.

[11] suggested a method for preservation of user privacy on the central server by adding uncertainty to the data. Before transferring personal data to the server, each user first "disguises" using a randomized perturbation technique. Therefore, the server (and also the attacker) can not find out the actual contents of users' profile. Although this method changes the user's original data, experiments show that the modified data still allows providing relatively accurate recommendations. This approach enhances users' privacy, but the users still depend on centralized, domain-specific servers that they subscribe to, on getting a recommendation. The capability to dynamically receive "on-the-fly" recommendations is limited.

In general, storing users' profiles on several locations reduces the risk of having the data exposed to an attacker in comparison to a storage on a single server. CF over a distributed setting of data repositories was initially proposed in [16]. This work presented a Peer-to-Peer architecture supporting product recommendations for mobile customers represented by software agents. The communication between the deployed agents used expensive routing mechanism based on network flooding that increased the communication overhead. An improved mechanism was proposed in [10], however it reduced the efficiency of neighborhood formation phase. The work in [14] elaborated the discussion on distributed CF. It developed a detailed taxonomy of distributed CF in recommender systems and presented different implementation frameworks for different domains of Electronic Commerce. PocketLens project [9] implemented and compared five distributed architectures for CF. It was found that no architecture is perfect, but the performance of content-addressable mechanism [13] is close to the performance of centralized CF algorithm.

Other approach to distributed recommendation is by completely eliminating the use of servers. A user creates a query by sending a part of his profile and a request for recommendation on this specific item. Other users decide to respond and send their information to the requester. However, this approach still requires transferring the users' profiles over the network, thus posing privacy issues. Two schemes for privacy-preserving CF were proposed in [2] and [3]. In these schemes the users control all of their own private data, while a community of users can compute a public "aggregate" of their data without disclosing individual users' data. The aggregate allows personalized recommendations to be computed by members of the community, or by outsiders. This approach protects users privacy in a distributed setting, but requires users group formation and complicated communication which are limitations in today's evolving dynamic environments

### 3 On the-fly recommendation generation

Users are the owners of their personal information. Thus, they should decide if, when and how to reveal parts of their user-profile. Recommendations may be requested in specific context – specific domain, in specific time and location. The users need to have continuous access to recommendation systems, not limited by availability of servers or predefined users groups.

The easiest way to support this request is by applying “Peer-to-Peer” recommendations (figure 1). Revealing a user profile is required in two different cases. The first is when a user requests a recommendation from other users. In this case the user must reveal his/her own profile in order to receive relevant information (recommendation). The second case is when a user decides to provide recommendation to other users by revealing his/her user profile, so the recommendation requester can use it in order to construct a neighborhood of similar users and generate a recommendation.




Fig. 1. Centralized vs. Decentralized Storage of Users' Profiles

In both cases users' privacy is at risk. We try to preserve user's privacy in the following way. Like Canny [2],[3], we believe that recommendations should be provided by individuals, at will. However, today's dynamic online environment prevents formation of communities and aggregation of users' profiles. For requesting recommendation, user may reveal only the relevant part of his profile, and even this part can be partially revealed. This solution is similar to the approach taken by Polat and Du in [11] and [12]. Only a subset of the features ratings should be provided. In addition, when providing a recommendation to another user, only a recommendation with a degree of similarity can be provided, instead of the actual user profile.

There are few questions that are posed when considering this approach. The main question is what portion of a user profile is needed in order to generate a recommendation. This is relevant for both sides – the recommendation requester and the recommendation provider. Another question is the question of trust – how trustworthy are the recommenders.

The current work addresses the first question and demonstrates that it is possible to use a relatively small portion of user profile in order to generate good recommendation. This is true for both the information requester and the responding users. The practical meaning is that users may protect their privacy simply by revealing small and partial portions of their profile when requesting information or providing a recommendation.

Though this approach improves the privacy of the responding users, it still allows to reveal the users' profile through a systematic attack using similarity requests. We

further increase the users’ privacy by obfuscating parts of their profiles in the process of calculating the similarity. Thus, the values collected by an attacker will not be reliable enough for a single user and complicate the task of inferring about the real contents user profile. However, since the system has many users such local rating obfuscation should not hamper too much the overall system performances.

## 4 Experimental results and discussion

In the experimental part of our work we used Jester dataset of jokes. Jester is a web-based joke recommendation system, developed at Berkeley University [5]. The database contains 4.1 million continuous ratings (-10.00 to +10.00) of 100 jokes from 73,421 users. Significant part of the users end up reading and rating all the jokes, so Jester dataset is relatively dense. In total, almost 50% of all possible ratings in the matrix are present.

We chose a subset of 1,024 users that rated all 100 jokes to get a dense matrix where every rating in the matrix corresponds to an actual user rating. We simulated decentralized distributed environment by a Java multi-threaded implementation. Each request was transferred to the relevant users, each user computed the similarity locally (possibly on an obfuscated profile), and returned the similarity rate and the required rating to the originator of the request. Upon receiving the responses from the other users, the originator generates the prediction locally as a weighted average of the ratings of the most similar users. Hence, the process for generating the prediction is performed in the same manner as in a centralized CF, except the similarity computation that is done distributively.

The first experiment aims to test how storing profiles at the client-side influences the accuracy of the generated recommendations. To measure the accuracy of the prediction we used Mean Average Error (MAE) [8] metrics. MAE was computed by:

$$MAE = \frac{\sum_{i=1}^N |p_i - r_i|}{N},$$

where  $N$  denotes the total number of the generated predictions,  $p_i$  is the  $i^{th}$  prediction, and  $r_i$  is the real  $i^{th}$  rating.

The MAE results obtained in the experiment are relatively low, between 0.15 and 0.18, implying that generated predictions are similar to the real ratings of the originating users. These MAE values are also similar to ones obtained at previous works that used the Jester dataset (initially presented in [5], and recently compared in [3]). Despite the fact that a centralized storage was distributed between the users, the actual ratings in the profiles remained unchanged. Thus, there is no reason to expect different MAE values.

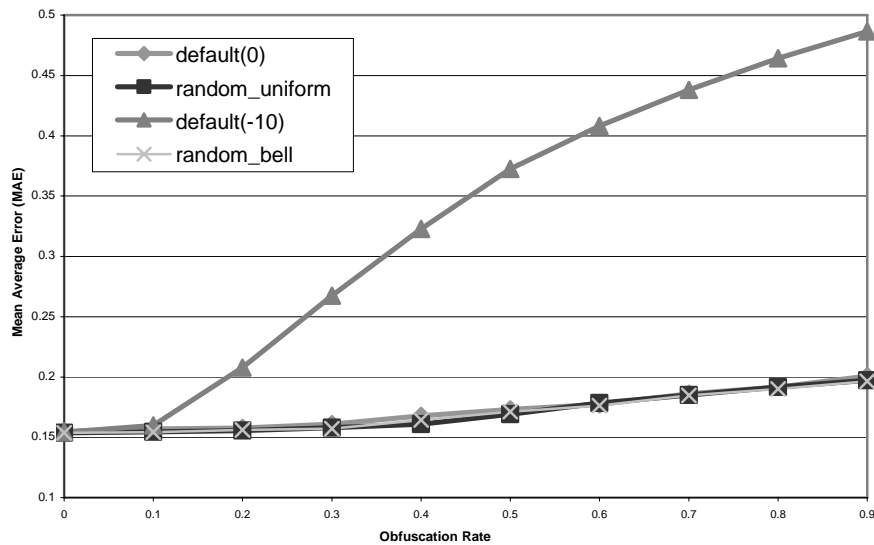
The second experiment aims at determining the influence of data obfuscation on the recommendation accuracy. We compared four different methods for modifying the data in users’ profiles. We measured the effect of gradually replacing increasing parts of the users’ profiles with either a predefined value or randomly chosen value. When replacing real values with predefined ones, we also tested the effect of the choosing various values.

Thus, we define three basic policies for modifying the data in users' profiles:

- Uniform Random obfuscation – real ratings values in the user's profile are substituted by a random values chosen uniformly in the scope of possible ratings (-10.00 to +10.00).
- Bell curved Random obfuscation – real ratings values in the user's profile are substituted by random values chosen using a bell curve distribution with similar statistics to the dataset.
- Default obfuscation( $x$ ) – real ratings values in the profile are substituted by a predefined value  $x$ .

To check the influence of  $x$  in obfuscation( $x$ ) policy on the accuracy of recommendation, we conducted this experiment with two different values of:  $x=0$  (which is close to the average of the ratings of the dataset), and  $x=-10$  (an extremely negative rating).

At each experiment we gradually increased the percentage of user profile that is modified (further referred as obfuscation rate) from 0.0 (the original profile is unchanged) to 0.9 (90% of the ratings in a profile of each user are modified). We produced a fixed testing set of 10,000 random jokes, and for each possible obfuscation rate we measured the MAE for the whole testing set. Figure 2 illustrates MAE results as a function of the obfuscation rate.



**Fig. 2.** MAE vs. Obfuscation Rate

Figure 2 shows that the performance of both Random policies and Default(0) obfuscation policies is similar. These policies do not drastically change the accuracy of the generated recommendations. The MAE rate slightly increases as the obfuscation rate increases, however the change is minor (from 0.15 to 0.2) and the prediction is still accurate. We explained it by considering that both the average rating value of the Default(0) obfuscation and the expectation rating value of the Uniform Random obfuscation is equal to 0. The expectation rating value of the bell curved Random ob-

fuscation is equal to  $I$ . These values are close to the average rating value of the ratings in the dataset. Thus, substituting the actual ratings with similar ratings creates only a small overall impact on the MAE computed over many users.

On the other hand, using the Default(-10) obfuscation policy, the actual ratings are substituted by a highly dissimilar value, (as it is far from the average value in the data set). As a result, the MAE rate increases linearly starting from 10% obfuscation rate.

#### 4.1 Discussion

The experimental results demonstrate that it is possible to use a relatively small portion of user profile in order to generate good recommendation. However, these results raise a number of interesting research questions related not only to privacy, but to other fundamental CF topics in general. In the rest of this section we will briefly describe these questions:

- The experiments were performed on a dense subset of Jester dataset. Despite that, we claim that obfuscating part of the data has the same impact as working with sparse dataset.
  - Will the results change when conducting the same experiments on a sparse dataset, e.g., MovieLens?
  - Could a differential obfuscation policy be developed to minimize loss of real data in sparse datasets?
- Analyzing the nature of the data. In Jester dataset (and supposedly in other datasets) there are many items that most users agree on their rating,
  - Can random selection provide similar results to a sophisticated recommendation system?
- The obfuscated results show that all the users have basically very similar profiles and still the prediction is good. This means that most users tend to prefer the same jokes.
  - Could it be confirmed on another data set?
  - What about “non-standard” users that have different preference, how will obfuscation influence them?
- Standalone attacks of malicious users through changing their ratings can not affect the accuracy of the global predictions.
  - How will our approach scale under an organized attack of multiple users hamper the functionality of CF?
- Our approach still requires transferring the originator profile over the network posing privacy issues.
  - Can the originator’s profile be perturbed in a similar way?
- In a real situation only a fraction of the peers are likely to respond.
  - How many peers will be needed to communicate with and ultimately what is the scalability and cost of the algorithm?

We believe that these questions require additional research before practical conclusions should be drawn regarding the contribution of the demonstrated approach.

## 5 Conclusions and Future Research

The need to protect of users' privacy triggers growing research efforts. Many eavesdroppers on the Web violate users' privacy and as a result users concerned about their privacy refrain from using useful Web services to prevent exposure. Additionally, in today's dynamic environments, formation of a community of users poses limitation on possibilities of information sharing. Users looking for personal information in various domains and situations may need to interact with different set of users every time.

This work proposes a simple and effective solution for preservation of the users' privacy during information sharing interaction. It employs a privacy-enhanced CF algorithm that allows creating dynamic and distributed recommendations. These recommendations are generated "on-the-fly" by letting the individual users participate in a virtual, distributed CF system. The users control when and what is the personal information they reveal.

Our approach stores users' profiles on several different locations and thus has the advantage of reducing the risk of having the users' data exposed to a malicious attacker. Moreover it can decrease the likelihood that the information could be collected for the purpose of transferring or selling and then be used in malicious way.

In order to further increase users' privacy, parts of the users' profiles are obfuscated in the process of calculating their similarity to the originator user. Thus, values collected by an attacker will not be reliable enough for a single user and complicate the task of inferring about the real contents of the user profile. As the system has many users such local obfuscations does not hamper the overall system performances.

### 5.1 Future Research

The current work demonstrated that it is possible to use a relatively small portion of user profile in order to generate good recommendation. However, in a real situation, only a fraction of the users are likely to respond. That means the user originating the request will have to send their data to a much larger set of peers. A natural extension of our approach would be to study the accuracy of privacy enhanced CF as a function of the number of peers who responded. This will answer an important question for determining how many peers to should be communicated with, and ultimately the scalability and cost of the algorithm.

Another future research direction is the problem extreme sparseness in the CF domains. Available movies databases (who have several thousand titles) are still at the low end of the number of choices, and are relatively dense. However, there are many more CDs or books, or TV shows to choose from (and in practice, these items follow a Zipf distribution). Thus, sparseness is an inescapable reality for most practical CF domains. Unfortunately, statistical obfuscation and sparseness do not correlate well together. Perturbing missing data items could swamp the information found in the real data. We plan to investigate how our approach behaves on real sparse datasets and possible ways to improve it.

## References

- [1] S.Brier, "How to Keep your Privacy: Battle Lines Get Clearer", The New York Times, 13-Jan-97.
- [2] J.Canny, "Collaborative Filtering with Privacy", in IEEE Symposium on Security and Privacy, Oakland, CA, 2002.
- [3] J.Canny, "Collaborative Filtering with Privacy via Factor Analysis", in proceedings of International ACM SIGIR Conference on Research and Development in Information Retrieval, Tampere, Finland, 2002.
- [4] L.F.Cranor, J.Reagle, M.S.Ackerman, "Beyond Concern: Understanding Net Users' Attitudes about Online Privacy", Technical report, AT&T Labs-Research, April 1999.
- [5] K.Goldberg, T.Roeder, D.Gupta, C.Perkins, "Eigentaste: A Constant Time Collaborative Filtering Algorithm", in Information Retrieval, 4(2), pp.133-151, 2001.
- [6] P.Harris, "It is Time for Rules in Wonderland", Businessweek 20, 2000.
- [7] J.Herlocker, J.A.Konstan, J.Riedl, "Explaining Collaborative Filtering Recommendations", in proceedings of ACM Conference on Computer Supported Cooperative Work, Philadelphia, PA, 2000.
- [8] J.L.Herlocker, J.A.Konstan, L.G.Terveen, J.T.Riedl, "Evaluating Collaborative Filtering Recommender Systems", in ACM Transactions on Information Systems, vol.22(1), pp.5-53, 2004.
- [9] B.N.Miller, J.A.Konstan, J.Riedl, "PocketLens: Toward a Personal Recommender System", in ACM Transactions on Information Systems, vol.22 (3), 2004.
- [10] T.Olsson, "Decentralised Social Filtering based on Trust", in proceedings of AAAI-98 Recommender Systems Workshop, Madison, WI, 1998.
- [11] H.Polat, W.Du, "Privacy-Preserving Collaborative Filtering Using Randomized Perturbation Techniques", in proceedings of International Conference on Data Mining, Melbourne, FL, 2003.
- [12] H.Polat, W.Du, "SVD-based Collaborative Filtering with Privacy", in proceedings of ACM Symposium on Applied Computing, Nicosia, Cyprus, 2004.
- [13] S.Ratnasamy, P.Francis, M.Handley, R.Karp, S.Shenker, "A Scalable Content-Addressable Network", in proceedings of ACM SIGCOMM, San Diego, CA, 2001.
- [14] B.M.Sarwar, J.A.Konstan, J.Riedl, "Distributed Recommender Systems: New Opportunities for Internet Commerce", a chapter in "Internet Commerce and Software Agents: Cases, Technologies and Opportunities", Idea Group Publishers, 2001.
- [15] J.B.Schafer, J.A.Konstan, J.Riedl, "E-Commerce Recommendation Applications", in Journal of Data Mining and Knowledge Discovery, vol. 5 (1/2), pp. 115-152, 2001.
- [16] A.Tveit, "Peer-to-Peer Based Recommendations for Mobile Commerce", in proceedings of the 1st International Workshop on Mobile Commerce, Rome, Italy, 2001.



# Privacy, Shilling, and The Value of Information in Recommender Systems

Shyong K Lam and John Riedl

GroupLens Research  
Computer Science and Engineering  
University of Minnesota  
Minneapolis, MN 55455  
{lam,riedl}@cs.umn.edu

**Abstract.** Recommender systems are an increasingly popular tool used by many consumers to help deal with information overload in today's marketplace. At the cost of some personal information, these systems are able to personalize a user's online experience and guide them toward making better decisions. This paper examines two issues relating to privacy in recommender systems: the value of information and shilling. The paper considers the privacy cost of recommender systems and proposes ways that the loss of privacy can be limited and balanced against recommendation and personalization quality.

## 1 Introduction

Consumers in today's marketplace are often overwhelmed with the number of different options available to them. To combat this problem, which is often called *information overload*, many people have turned to recommender systems. These are tools that use opinions about items in some information domain in order to make recommendations to a user regarding which items she may wish to consider. One example of a recommender system is *MovieLens* (<http://www.movielens.org>), which is a system that makes personalized recommendations suggesting movies that a user might like based on the movies she has seen and has expressed an opinion about. Using recommender systems, online retailers can personalize their virtual storefronts specifically for each consumer to help them find and buy things.

Unfortunately, while recommenders provide benefit to users as a decision-making tool, this benefit often comes with a cost in privacy. Personalized recommenders like *MovieLens* require some information about the user's taste before they are able to make personalized recommendations. (Some recommenders, such as *Amazon.com*'s customer reviews, provide recommendations with no information about the user's tastes, but such systems are not truly personalized — everybody sees the same reviews.)

In our work, we largely focus on personalized recommender systems. In particular, we are interested in ones that use automated collaborative filtering (ACF), which refers to algorithms that generate recommendations on the

basis that people who have expressed similar opinions in the past are likely to share opinions in the future. Two commonly used ACF algorithms are user-based k-Nearest-Neighbor[1] and item-based k-Nearest-Neighbor[2]. These algorithms are commonly classified as *memory-based* algorithms where the entire set of preference information is used in producing recommendations.

A second class of ACF algorithms consists of *model-based* algorithms in which the preference information is processed into a model that can in turn be used to compute recommendations. An example of such an algorithm is the one based on Bayesian networks as described by Breese et al. [3]. These models contain the preference information in a reduced form, and it is difficult – and perhaps impossible – to recover the original preference information given just the model. Using a model-based algorithm allows the recommender system to preserve its users’ privacy in that it can distribute the models to its users who can in turn compute recommendations without directly providing preference information to the server. Model-based algorithms of this type make possible the separation of data about *self* from data about *others*. In principle, data about self might be maintained on the client, and not shared with the centralized server. However, most model-building algorithms retain the requirement that sufficiently many users must provide the system with preference information in order to have a high-quality model built in the first place.

There has been substantial theoretical work done in studying this and other privacy issues in ACF-based recommender systems. Ramakrishnan et al [4] describe a graph-theoretic model showing how information can be inferred about certain classes of users by observing the recommendations made by a system. Canny [5] and Miller et al. [6] suggest new privacy-preserving paradigms for recommender systems that use peer-to-peer and encryption to prevent the system operator from learning any explicit preference information about its users, and to limit what users can learn about each other. Even with these restrictions on information flow, the proposed systems are able to produce recommendations with accuracy comparable to existing non-privacy-preserving systems.

We seek to extend this work by exploring steps to promote privacy that can be taken with currently-deployed systems that may have already collected a significant amount of data about their users. Most of the discussion applies equally to memory-based or model-based algorithms. The theme of this paper is to informally explore the costs and benefits of limiting the amount of information kept by a recommender about a user. By using information-theoretic techniques, we believe it may be possible to selectively decide which information about a user to retain and which to discard while being able to perform a reasonable amount of personalization. In doing this, however, it is possible that the recommender system becomes more vulnerable to malicious attacks by third parties.

## 2 Value of Information

A personalized recommendation algorithm typically requires input from the user population in order to make recommendations. In general, the more information

that is known about the users, the more accurate the recommendations are. This presents a conundrum when considering privacy in systems that use such algorithms — more data leads to better personalization, but acquiring data can be invasive to user privacy. Ideally, one would like to find a balance where the system is able to make good recommendations while not requiring users to give up too much information about themselves.

The accuracy of an algorithm with respect to the amount of information known about the user follows a diminishing returns curve. That is, once a certain amount is known about a user, obtaining further information is only marginally useful. This raises the possibility of finding a “sweet spot” that maximizes the recommendation accuracy per unit of information known about the user.

Some people, particularly advertisers, seek to provide personalization based on a small amount of information. For instance, recommendations might be based on demographic data (e.g. the usual “ZAG” information — zip code, age, gender), or generalized preferences of attributes describing the items involved (in movies, this might mean the user’s favorite genres). These personalization methods only have a modest effect on privacy as the information given up by the user cannot easily reveal the user’s identity.

In contrast, highly personalized recommenders, such as those based on ACF, require a high degree of personal preference information from the user. These requirements lead to larger privacy concerns since this level of preference information may reveal substantial personal information about the user. There are many open privacy-related research questions around the elicitation of such preference information, including: (a) How much information is given up by a user when she provides an opinion about an item? (b) Does the amount of information content vary by item? (c) How does the privacy lost relate to the information gained by the recommender system?

Note that the answers to these questions are likely to vary by domain. In some domains, such as music CDs, users may be relatively open to sharing their tastes with others. In other domains, such as medical information, users may have serious concerns about sharing their preferences with anyone, because of the potential harm should the information leak to colleagues or potential employers. In still other domains, such as scientific research papers, the sensitivity of the information may vary with time. While working on a paper, a researcher may not want others to know what related work she is studying; once the paper is published, the list of references is publicly available and no longer presents a privacy concern. In the discussion we try to focus on aspects of privacy in recommenders that are likely to be domain-independent.

Intuitively, the goal of preference information is to differentiate a user from her peers. Preference information that does a better job of differentiating among users should be inherently more useful in personalizing a system for that user. For instance, knowing that a user likes the movie “Toy Story” reveals less about her than knowing that she likes “Fahrenheit 9/11.” The former is a universally-liked movie, while the latter is more polarizing with a higher level of disparity among users’ opinions. This is the basis of an idea proposed by Pennock and

Horvitz that says if one can calculate how useful a given piece of information is, that is, if one had a “value-of-information” metric, then one can tune a system to optimize its data collection process by choosing to solicit user preferences on items that carry the most value [7].

In past work we have explored the related issue of eliciting information from new users in a recommender in ways that optimize both the required user effort and initial recommendation accuracy [8, 9]. Using the value of information (VOI) concept, we were able to build VOI-aware MovieLens interfaces that reduced the user effort needed to start receiving recommendations. Moreover, the accuracy of those recommendations was improved compared to ones made based on the initial user models built for users that did not use the VOI-aware interface.

In the interest of user privacy, this kind of approach may be comforting to some users in that fewer discrete pieces of information (e.g. movie ratings) need to be provided before the system becomes accurate. However, the information theory involved says that the user has given up an equivalent amount of information about herself as she would have with an unoptimized approach. On the other hand, it is possible to use VOI to measure how much information is needed to make good recommendations, and then to stop collecting new information from the user once that point is reached. More generally, a recommender system can use VOI to bound the amount of information collected about a user to some optimal level with respect to both privacy and recommendation quality.

These ideas lead to the following set of questions about the role of VOI in preserving user privacy for existing users in a recommender system that may comprise an interesting follow-up program of research:

## 2.1 Amount of Data

How much data is needed from a user to make good recommendations? How do we calculate the “sweet spot”? It is desirable for the system to know this so that it could stop eliciting data from a user when it feels it has built a sufficiently good user model. A related issue here is how this may change over time. As the system or the user evolves, or as new items are introduced, will more information be needed to maintain high-quality recommendations?

One challenge with limiting the amount of data provided by a user is that the system might then recommend items the user already possesses. For instance, a recommender for music CDs might prefer to filter out from recommendation lists CDs the user already owns. To support this feature, the system might encourage the user to upload complete information about her collection. The user would then have to balance her preference for privacy with her preference for personally filtered recommendations. One privacy-preserving solution might be to have a client-side filter select from among the recommended items those to display to the user. The centralized system could then form recommendations based on the modest amount of data it needs for quality recommendations.

## 2.2 User Interface

What should the user interface look like, especially after the system thinks it has learned enough about the user? What if the user *wants* to tell us more about herself? How does one present a privacy-preserving recommender system in an understandable way? In our experiences with MovieLens, we have found no shortage of people willing to provide hundreds and sometimes thousands of movie ratings. Indeed, user feedback from our periodic surveys reveals that rating movies is among the leading reasons people have for using the system! These observations that some users are perfectly willing to give up their privacy may make it tricky to create a usable interface that effectively conveys the privacy-preserving aspects of the recommender system.

## 2.3 Selectively Discarding Data

If we find ourselves knowing “too much” about a user, which particular pieces of data should be kept? This is strongly tied to VOI, but the answer to this question is not necessarily the information with the highest value. If “low-valued” information is discarded from many users’ models, then perhaps that information is no longer low-valued since it has become more rare. Choosing an appropriate set of data that balances both the benefit to the overall system and the quality of each individual user model seems to be a challenging task.

## 2.4 Impact on CF Algorithms

How well do current collaborative filtering algorithms operate in reduced-data environments? Clearly, both recommendation accuracy and coverage might be affected. Some algorithms such as SVD seem more naturally suited for sparse data sets [10] — are they even better if given selectively chosen data? Is the sparse data inherently less noisy, and if so, could it even lead to *better* recommendations using specialized algorithms? Another effect might be that the algorithm is more susceptible to shilling attacks. This will be described further in section 3.

# 3 Shilling

One of the primary uses for a recommender system is to help people make decisions, whether it be which movie to see, which restaurant to eat at, or what web site to visit. Naturally, this makes recommender systems very interesting to people with vested interests in what people choose to consume. For instance, a restaurant owner would be more successful if more people ate at his establishment, so it is within his best interests to have it recommended often by a restaurant recommender system. One way to do this is to provide good service to garner a good reputation among restaurant diners. In turn, this would likely lead to more frequent recommendation as users express high opinions of the restaurant on the recommender system.

However, a more underhanded and often cheaper way to increase recommendation frequency is to manipulate or trick the system into doing so. This can be done by having a group of users (human or agent) use the recommender system and provide specially crafted “opinions” that cause it to make the desired recommendation more often. Instances of these manipulations have been observed in the past. For example, it has been shown that a number of book reviews published on Amazon.com are actually written by the author of the book being reviewed<sup>1</sup>. A consumer trying to decide which book to purchase could be misled by such reviews into believing that the book is better than it really is. The privacy Amazon provides to its reviewers allows this to happen, as reviews are usually only attributed to a pseudonym. Unfortunately, there is no clear-cut solution here since reviewers often wish to maintain their privacy, which can conflict with the reader’s desire to be able to know how much they can believe a review. As a step in solving this problem, Amazon has introduced a “Real Name” feature to their review system that specially identifies reviewers who have elected to divulge their “true” identity, and thus might be more trustworthy.

Our previous work [11] has shown that these types of “shilling” attacks are indeed effective and are relatively easy to carry out with collaborative filtering algorithms in use today. An attacker only needs to know a small amount of information about the user and item population in order to perform an attack that increases the number of times some particular item (or set of items) is recommended. Furthermore, the attacks are non-trivial to detect with typical measures of recommender system performance.

A concern of introducing more privacy to a recommender system is whether it might make these attacks easier or more common, as the pseudonymity or anonymity provided by privacy usually invites abuse of this nature (consider what usually happens on any web-based forum or bulletin board when no controls are placed on posting messages). In particular, the issue of privacy in recommender systems raises the following questions regarding shilling attacks.

### 3.1 Attack Effectiveness

Do shilling attacks become more effective against privacy-preserving recommender systems? As additional privacy is introduced to a recommender system, the opportunities for attacks can increase considerably. Our work [11] shows that attacks that target recommendation frequency of low-information items (i.e. ones with few ratings) are more effective than attacks against high-information items. In a system that tries to maintain a minimal amount of information about its members, it is possible that *every* item might have sufficiently few ratings to be vulnerable to highly-effective attacks.

### 3.2 Attack Difficulty

Are shilling attacks more or less difficult to mount against privacy-preserving recommender systems? As mentioned above, more individual items might be

---

<sup>1</sup> <http://www.onlinesecurity.com/links/links837.php>

come ideal targets for effective attacks. On the other hand, if the recommender system only keeps some subset of data provided by each user, an attack strategy will need to take that into consideration, both for the users being targeted and for the “false” users introduced by the attack. This would likely require the attacker to know more about the system being attacked, thus increasing the cost of an attack.

Another possible impeding factor in an attack is the interface presented to users. A VOI-aware interface such as the ones used in our past work [8, 9] can control which items may be rated by a user in order to maximize the information gain per collected rating. This significantly constrains what an attacker can do and could make it more difficult to impact the system in precise ways.

### 3.3 Attack Detection

In a privacy-preserving recommender system, is it easier or harder to detect an attack? One might theorize that in a low-data environment, it becomes easier to identify atypical patterns that are indicative of an attack. If found to be true, this would certainly be a boon to recommender system operators. On the other hand, discarding some of the data entered by a shilling agent might leave the remaining data looking “more human,” and hence harder to detect.

## 4 Conclusion

The issue of privacy in recommender systems is a rich area with many aspects that have yet to be explored. Recommenders – especially highly personalized recommenders like those using automated collaborative filtering – raise important issues about how the data they collect impinges on user privacy.

In this paper we explored two aspects of recommender systems that relate to these important privacy questions: VOI and shilling. Previous work on VOI shows that it can be used to more effectively collect information from new users. We believe it can similarly be used to determine when to stop collecting information to properly balance the privacy given up by users with the quality of the recommendations, and to intelligently choose which information to discard if “too much” is known about a user. The challenge of shilling is that the aforementioned privacy protections may make shilling easier, especially if they reduce the amount of information the recommender system keeps about each user.

This is, however, just the tip of the iceberg. There are many other aspects of privacy in recommender systems, including its role in community-oriented systems where participants interact with each other on a regular basis, or in small-world systems in which most or all of the participants know each other. Here, privacy might not be as big a concern to users as it might be with, say, a large corporate-owned system. In fact, users might even want ways to share validated personal information, so they know something about who they are interacting with. Successful systems will likely have layers of privacy, with users in charge of what they see from others and what they show to others.

## 5 Acknowledgments

Thanks to the workshop referees for their helpful suggestions, and to members of GroupLens Research at the University of Minnesota for many fruitful discussions. Particular thanks are due to our colleagues Al Mamunur Rashid, Istvan Albert, Dan Cosley, Sean M. McNee, and Joseph Konstan, our co-authors on the VOI research [8, 9]. This work was supported by grants from the NSF (DGE 95-54517, IIS 96-13960, IIS 97-34442, IIS 99-78717, and IIS 01-02229).

## References

1. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., Riedl, J.: GroupLens: an open architecture for collaborative filtering of netnews. In: CSCW '94: Proceedings of the 1994 ACM conference on Computer supported cooperative work, Chapel Hill, North Carolina, United States, ACM Press (1994) 175–186
2. Sarwar, B., Karypis, G., Konstan, J., Reidl, J.: Item-based collaborative filtering recommendation algorithms. In: WWW '01: Proceedings of the tenth international conference on World Wide Web, Hong Kong, ACM Press (2001) 285–295
3. Breese, J.S., Heckerman, D., Kadie, C.: Empirical analysis of predictive algorithms for collaborative filtering. In: Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI-98). (1998) 43–52
4. Ramakrishnan, N., Keller, B.J., Mirza, B.J., Grama, A., Karypis, G.: Privacy risks in recommender systems. *IEEE Internet Computing* **5** (2001) 54–62
5. Canny, J.: Collaborative filtering with privacy via factor analysis. In: SIGIR '02: Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval, Tampere, Finland, ACM Press (2002) 238–245
6. Miller, B.N., Konstan, J.A., Riedl, J.: PocketLens: Toward a personal recommender system. *ACM Transactions on Information Systems* **22** (2004) 437–476
7. Pennock, D.M., Horvitz, E., Lawrence, S., Giles, C.L.: Collaborative filtering by personality diagnosis: A hybrid memory and model-based approach. In: UAI '00: Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence, Stanford, CA, Morgan Kaufmann Publishers Inc. (2000) 473–480
8. Rashid, A.M., Albert, I., Cosley, D., Lam, S.K., McNee, S., Konstan, J.A., Riedl, J.: Getting to know you: Learning new user preferences in recommender systems. In: Proceedings of the 2002 International Conference on Intelligent User Interfaces, San Francisco, CA (2002) 127–134
9. McNee, S.M., Lam, S.K., Konstan, J.A., Riedl, J.: Interfaces for eliciting new user preferences in recommender systems. In: User Modeling, Johnstown, PA, USA, Springer Verlag (2003) 178–187
10. Sarwar, B.M., Karypis, G., Konstan, J.A., Riedl, J.: Application of dimensionality reduction in recommender system – a case study. In: ACM WebKDD 2000 Web Mining for E-Commerce Workshop, Boston, MA, USA (2000)
11. Lam, S.K., Riedl, J.: Shilling recommender systems for fun and profit. In: WWW '04: Proceedings of the 13th international conference on World Wide Web, New York, NY, USA, ACM Press (2004) 393–402

## Author Index

Ajay Brar 47  
Alfred Kobsa 35  
Andreas Pfitzmann 67  
Boris de Ruyter 15  
Brian Richardson 55  
Chris C. Demchak 1  
Elke Franz 67  
Evelien Perik 15  
Francesco Ricci 75  
Hagen Währig 67  
Hilko Donker 67  
Jim Greer 55  
John Riedl 84  
Judy Kay 47  
Katja Liesebach 67  
Katrín Borcea 67  
Kurt D. Fenstermacher 1  
Sarah Spiekermann 3  
Panos Markopoulos 15  
Peter de Vries 23  
Shlomo Berkovsky 75  
Shyong K Lam 84  
Thea van der Geest 23  
Tsvi Kuflik1 75  
Willem Pieterse 23  
Yang Wang 35  
Yaniv Eytan 75